

JC921 U.S. PTO  
08/31/00

9-5-00 A

Please type a plus sign (+) Inside this box +

<b>UTILITY PATENT APPLICATION TRANSMITTAL</b> <small>(Only for new nonprovisional applications under 37 C.F.R. § 1.53(b))</small>		Attorney Docket No. <b>15311-2284</b>
		First Inventor or Application Identifier <b>David Golden et al.</b>
		Title <b>PARTITION CONFIGURATION VIA SEPARATE MICROPROCESSORS</b>
		Express Mail Label No. <b>EL705595084US</b>

<b>APPLICATION ELEMENTS</b> See MPEP chapter 600 concerning utility application contents		Assistant Commissioner for Patents Box Patent Application Washington, DC 20231
---	--	--

1. <input checked="" type="checkbox"/> *Fee Transmittal Form (e.g., PTO/SB/17) <i>(Submit an original and a duplicate for fee processing)</i>	6. <input type="checkbox"/> Microfiche Computer Program (Appendix)
2. <input checked="" type="checkbox"/> Specification [Total Pages <b>107</b> ] <i>(preferred arrangement set forth below)</i>	7. Nucleotide and/or Amino Acid Sequence Submission <i>((if applicable, all necessary))</i>
- Descriptive title of the Invention	a. <input type="checkbox"/> Computer Readable Copy
- Cross References to Related Applications	b. <input type="checkbox"/> Paper Copy (Identical to computer copy)
- Statement Regarding Fed sponsored R & D	c. <input type="checkbox"/> Statement verifying identity of above copies
- Reference to Microfiche Appendix	
- Background of the Invention	
- Brief Summary of the Invention	
- Brief Description of the Drawings <i>(if filed)</i>	
- Detailed Description	
- Claim(s)	
- Abstract of the Disclosure	
3. <input checked="" type="checkbox"/> Drawing(s) [Total Sheets <b>10</b> ]	<b>ACCOMPANYING APPLICATION PARTS</b>
4. Oath or Declaration [Total Pages <b>  </b> ]	8. <input type="checkbox"/> Assignment Papers (cover sheet & document(s))
a. <input type="checkbox"/> Newly executed (original copy)	9. <input type="checkbox"/> 37 C.F.R. §3.73(b) Statement <i>(when there is</i> <input type="checkbox"/> Power of Attorney <i>an assignee)</i>
b. <input type="checkbox"/> Copy from a prior application (37 C.F.R. § 1.63(d)) <i>(for continuation/divisional with Box 17 completed)</i>	10. <input type="checkbox"/> English Translation Document <i>(if applicable)</i>
<b>[Note Box 5 below]</b>	
<b>DELETION OF INVENTOR(S)</b>	
Signed statement attached deleting inventor(s) named in the prior application, see 37 C.F.R. §§ 1.63(d)(2) and 1.33(b).	
5. <input type="checkbox"/> Incorporation By Reference <i>(useable if Box 4b is checked)</i>	12. <input type="checkbox"/> Preliminary Amendment
The entire disclosure of the prior application, from which a copy of the oath or declaration is supplied under Box 4b, is considered to be part of the disclosure of the accompanying application and is hereby incorporated by reference therein	
13. <input checked="" type="checkbox"/> Return Receipt Postcard (MPEP 503) <i>(Should be specifically itemized)</i>	
*Small Entity	
14. <input type="checkbox"/> Statement(s) <input type="checkbox"/> Statement filed in prior application, <i>((PTO/SB/09-12))</i> Status still proper and desired	
15. <input type="checkbox"/> Certified Copy of Priority Document(s) <i>(if foreign priority is claimed)</i>	
16. <input type="checkbox"/> Other:	

\*NOTE FOR ITEMS 1 & 14: IN ORDER TO BE ENTITLED TO PAY SMALL ENTITY FEES, A SMALL ENTITY STATEMENT IS REQUIRED (37 C. F. R. §1.27), EXCEPT IF ONE FILED IN A PRIOR APPLICATION IS RELIED UPON (37 C. F. R. § 1.28).

17. If a CONTINUING APPLICATION, check appropriate box and supply the requisite information below and in a preliminary amendment.

Continuation     Divisional     Continuation-in-part (CIP)    of prior application No.: /

Prior application Information: Examiner

Group/Art Unit:

**18. CORRESPONDENCE ADDRESS**

<input checked="" type="checkbox"/> Customer Number or Bar Code Label	<input type="checkbox"/> Correspondence address below				
(Insert Cus					
24267					
Name	A. Sidney Johnston			PATENT TRADEMARK OFFICE	
Address	Cesari and McKenna 88 Black Falcon Avenue				
City	Boston	State	MA	Zip Code	02210
Country	U. S. A.	Telephone	(617) 951-2500	Fax	(617) 951-3927

Name (Print/Type)	A. Sidney Johnston	Registration No. (Attorney/Agent)	29,548
Signature	<i>A. Sidney Johnston</i>		
Date	August 31, 2000		

**IN THE UNITED STATES PATENT AND TRADEMARK OFFICE**

In Re The Application of: )  
David Golden et al. )  
Serial No.: Not Yet Assigned ) Examiner: Not Yet Assigned  
Filed: August 31, 2000 )  
For: PARTITION CONFIGURATION ) Art Unit: Not Yet Assigned  
VIA SEPARATE MICROPROCES- )  
SORS )  
Cesari and McKenna, LLP  
88 Black Falcon Avenue  
Boston, MA 02210  
August 31, 2000

**EXPRESS-MAIL DEPOSIT**

“Express Mail” Mailing-Label Number: EL705595084US

The following papers are being deposited with the United States Postal Service  
“Express Mail Post Office to Addressee” service pursuant to 37 C.F.R. §1.10:

Patent Application ( 107 pages including 10 claims )  
 Utility Patent Application Transmittal Letter  Informal Drawings ( 110 sheets )  
 Check in the amount of \$690.00  Fee Transmittal letter  
 Unexecuted Declaration

UNITED STATES PATENT APPLICATION

*of*

5

**David Golden**

**Dennis Mazur**

*and*

**Richard Bracken**

*for a*

15

**PARTITION FORMATION USING MICROPROCESSORS IN A MULTIPROCESSOR  
COMPUTER SYSTEM**

20

## INCORPORATION BY REFERENCE OF RELATED APPLICATIONS

This patent application is related to the following co-pending, commonly owned U.S. Patent Applications, all of which were filed on even date with the within application for United States Patent, which are each hereby incorporated by reference in their entirety:

5       U.S. Patent Application Ser. No. (15311-2281) entitled ADAPTIVE DATA PREFETCH PREDICTION ALGORITHM;

          U.S. Patent Application Ser. No. (15311-2282) entitled UNIQUE METHOD OF REDUCING LOSSES IN CIRCUITS USING V<sup>2</sup> PWM CONTROL;

10      U.S. Patent Application Ser. No. (15311-2283) entitled IO SPEED AND LENGTH PROGRAMMABLE WITH BUS POPULATION;

          U.S. Patent Application Ser. No. (15311-2285) entitled SYSTEM AND METHOD FOR USING FUNCTION NUMBERS TO INCREASE THE COUNT OF OUTSTANDING SPLIT TRANSACTIONS;

15      U.S. Patent Application Ser. No. (15311-2286) entitled SYSTEM AND METHOD FOR PROVIDING FORWARD PROGRESS AND AVOIDING STARVATION AND LIVELOCK IN A MULTIPROCESSOR COMPUTER SYSTEM;

          U.S. Patent Application Ser. No. (15311-2287) entitled ONLINE ADD/REMOVAL OF SERVER MANAGEMENT INFRASTRUCTURE;

20      U.S. Patent Application Ser. No. (15311-2288) entitled AUTOMATED BACKPLANE CABLE CONNECTION IDENTIFICATION SYSTEM AND METHOD;

          U.S. Patent Application Ser. No. (15311-2289) entitled AUTOMATED BACKPLANE CABLE CONNECTION IDENTIFICATION SYSTEM AND METHOD;

          U.S. Patent Application Ser. No. (15311-2290) entitled CLOCK FORWARD INITIALIZATION AND RESET SIGNALING TECHNIQUE;

25      U.S. Patent Application Ser. No. (15311-2292) entitled PASSIVE RELEASE AVOIDANCE TECHNIQUE;

U.S. Patent Application Ser. No. (15311-2293) entitled COHERENT TRANSLATION  
LOOK-ASIDE BUFFER;

U.S. Patent Application Ser. No. (15311-2294) entitled DETERMINISTIC HARDWARE  
BEHAVIOR BETWEEN MULTIPLE ASYNCHRONOUS CLOCK DOMAINS THROUGH

- 5 THE NOVEL USE OF A PLL; and

U. S. Patent Application Ser. No. (15311-2306) entitled VIRTUAL TIME OF YEAR CLOCK.

## FIELD OF THE INVENTION

This invention relates to the control of computer systems, and more particularly to the control of  
5 multiprocessor systems.

## BACKGROUND

It is standard engineering practice to assemble a plurality of processors into multiprocessor system. It is also standard engineering practice to couple the plurality of processors by a bus for the processor to exchange control signals, where the bus is usually referred to as  
10 the control bus. The processors can then exchange control information useful to operation of operating systems executing on the processors. For example, in the event that one of the processors fails, its failure can be detected by the other processors through control message protocols such  
15 as keep alive messages, etc.

More complex multiprocessor systems use auxiliary microprocessors to gather information about the main processors, and to communicate with other microprocessors using the control bus. An example of the use of microprocessors communicating over a control bus for control of  
20 processors in a multiprocessor system is described in co-pending and commonly owned United States Patent Application Serial No. 09/545,781, filed on April 7, 2000.

A disadvantage of using a control bus to exchange control information between processors, either with or without microprocessors to gather and to transmit the control information, is  
25 that the arrangement is limited in flexibility and expandability. For example, the control bus is ordinarily part of a backplane, and the backplane is fixed by the hardware employed. Also the arrangement is limited in that the input/output structures of the multiprocessor system are often connected to main busses within the backplane, and it is difficult to individually control the input and output devices using signals transmitted over the control bus.

There is needed a flexible system for control of processors in a multiprocessor system, and where the control system may be conveniently expanded or contracted as processors are added to or removed from the multiprocessor system.

5

## SUMMARY OF THE INVENTION

The invention is a control system using microprocessors which communicate through a Local Area Network (private LAN) to control operation of both processors and input and output subsystems (IO system) of a multiprocessor computer system. The processors each have memory associated therewith, and each processor has an IO system comprising a plurality of busses such as PCI busses, associated therewith. The processors are cabled together in a mesh arrangement so that messages can be transferred between any of the processors and delivered to memory associated with the destination processor, or delivered to an IO system associated with the destination processor, etc.

15 The microprocessors are powered on when power is applied to the chassis of the multiprocessor system, and the microprocessors then control the processors of the multiprocessor system, including applying power to the processors, forming hard partitions containing selected processors, computing routes for messages from a processor to a memory associated with any processor for read and write transactions, computing routes for messages to IO subsystems associated with any processor of the hard partition, forming partition boundaries so that processors in one hard partition cannot read and write to memory or IO systems associated with processors in another hard partition, forming soft partitions of processors, controlling boot-up of operating systems executing on the processors of the multiprocessor computer system, removing power from a failed processor, providing power to a repaired processor, etc.

20 25 There is a microprocessor associated with each processor, and a microprocessor associated with each IO subsystem, and these microprocessors communicate through the private LAN. Each microprocessor maintains a data base giving the configuration of all of the processors in the multiprocessor computer system, and this data base is maintained current by being transferred through the private LAN as changes occur in the system. By use of this data base the microprocessors are enabled to perform their control functions. In a preferred embodiment of the inven-

tion, the processors of the multiprocessor computer system are mounted with two processors on one backplane, there are four (4) backplanes in a rack along with the memory associated with each processor and a drawer for each processor backplane containing the IO system associated with the two processors of that backplane, and in the rack there is one of the microprocessors controlling the eight (8) processors of the rack. The microprocessors communicate by an Ethernet LAN comprising twisted pair media coupling to an Ethernet hub.

Other and further aspects of the present invention will become apparent during the course of the following description and by reference to the accompanying drawings.

10

## BRIEF DESCRIPTION OF THE DRAWINGS

Referring now to the drawings, in which like numerals represent like parts in the several views:

Fig. 1 is a schematic block diagram of a Symmetrical Multiprocessor (SMP) system comprising a plurality of dual processor (2P) modules interconnected to form a two dimensional (2D)-torus mesh configuration;

Fig. 2 is a schematic block diagram of a 2P module of Fig. 1;

Fig. 3 is a schematic diagram of a memory subsystem of the SMP system;

Fig. 4 is a schematic block diagram showing the organization of the memory subsystem of Fig. 3;

Fig. 5 is a schematic block diagram of an IO7 of an I/O subsystem of the SMP system;

Fig. 6 is a schematic diagram of an illustrative embodiment of four (4) 8P drawers of the SMP system mounted within a standard 19-inch rack;

Fig. 7 is a schematic block diagram of an I/O drawer of the SMP system; and

Fig. 8 is a schematic block diagram of a server management platform for the SMP system.

25

Fig. 9 is a block diagram of a multi processor system showing interconnection of the system management components:

Fig. 10 is a block diagram of firmware used in system management components;

Fig. 11 is a table showing External Server Management Commands;

30

Fig. 12 is a table giving Internal Server Management Commands;

Fig. 13 is a table giving Time of Year Data, corresponding to BB\_WATCH;  
Fig. 14 is a table giving messages to provide compatibility to older systems;  
Fig. 15 is a block diagram giving Server Management Hardware Overviews;  
Fig. 16 is a block diagram giving MBM Hardware Overview;  
5 Fig. 17 is an isometric view of a Rack and Box with Thumbwheel switches;  
Fig. 18 is a table giving MBM task attributes;  
Fig. 19 is a block diagram giving an MBM firmware overview;  
Fig. 20 is a diagrammatic sketch showing group members electing a group leader;  
10 Fig. 21 is a table giving Powerup Flow with group relationships;  
Fig. 22 is a LAN versus field programmable gate array (FPGA) capability matrix;  
Fig. 23 is a PMU Server Block Diagram;  
Fig. 24 is a PMU Server Received Command Handling table;  
15 Fig. 25 is a table giving PMU Server Originating Commands;  
Fig. 26 is a diagrammatic sketch giving an exemplary system configuration;  
Fig. 27 is a table giving a Show Configuration Flow Diagram, Part 1;  
Fig. 28 is a table giving a Show Configuration Flow Diagram, Part 2;  
Fig. 29 is a Show Configuration Flow Diagram, Part 3;  
20 Fig. 30 is a table giving a Show Configuration Sample Output;  
Fig. 31 is a block diagram of a Partition Request Sources Process;  
Fig. 32 is a table giving a method of Partition Coordinator for Handling Requests;  
Fig. 33 is a table giving a Partition Coordinator Commands Issued Reference;  
Fig. 34 is a State Diagram of a Partition;  
25 Fig. 35 gives Inputs and Outputs of a Router Configuration Algorithm;  
Fig. 36 is a table giving a Routing Glossary;  
Fig. 37 is a block flow diagram of Partition Coordination;  
Fig. 38 is a block diagram of a process Creating a Hard Partition;  
Fig. 39 is a block diagram of a process Creating a Sub-Partition;  
Fig. 40 is a block diagram of a process creating a new partition flow diagram, Part 1;  
30 Fig. 41 is a table giving a process for creating a New Partition Flow Diagram, Part 2;  
Fig. 42 is a partition Start Flow Diagram, Reset state;

Fig. 43 is a partition Start Flow Diagram, Diagnostic State;  
Fig. 44 is a partition Start Flow Diagram, Configure Router;  
Fig. 45 is a partition Start Flow Diagram, Running;  
Fig. 46A is a Flow Diagram, “add EV7 Flow Diagram, Part 1”;  
Fig. 46B is a Flow Diagram, “Add EV7 Flow Diagram, Part 2”;  
Fig. 47 is a Flow Diagram for an Add versus Move;  
Fig. 48 is a block diagram illustrating Destroying a soft partition;  
Fig. 49 is a block diagram illustrating Destroying a hard partition;  
Fig. 50 is a flow diagram illustrating a “EV7 Failure/Replacement Flow Diagram,  
Part 1”;  
Fig. 51 is a flow diagram illustrating “EV7 Failure/Replacement Flow Diagram, Part 2”;  
Fig. 52 is a Set Membership Configuration Flow Diagram;  
Fig. 53 is an IP Cable Configuration Block Diagram;  
Fig. 54 is a table giving an EV7 Coordinate addressing relationship to thumbwheel ad-  
ressing;  
Fig. 55 is a IO7 Cabling Block Diagram;  
Fig. 56 is a get Cable Configuration Block Diagram;  
Fig. 57 is an IP Cable Addition/Deletion Block Diagram;  
Fig. 58 is a PMU Cabling Assistant Block Diagram;  
Fig. 59 is a Proxy Forwarding Block Diagram;  
Fig. 60 is a Virtual Console Terminal Overview;  
Fig. 61 is a table giving a Virtual Terminal Telnet port numbers;  
Fig. 62 is a table giving a Virtual Terminal Flow Diagram;  
Fig. 63 is a table giving a set base Time Flow Diagram;  
Fig. 64 is a Block Diagram giving a DHCP;  
Fig. 65 is a table giving a Flash Memory Layout;  
Fig. 66 is a table giving an Image Header;  
Fig. 67 is a table giving a SRM Environment Variables Flow Diagram;  
Fig. 68 is a Firmware Load and Upgrade Block Diagram;  
Fig. 69 is an Upgrading CMM Firmware Flow Diagram;

Fig. 70 is a table giving an Error Log Entry Format;  
Fig. 71 is a table giving an Error Entry Data Format;  
Fig. 72 is a diagram giving an Error Reporting Flow Diagram;  
Fig. 73 is a flow diagram for logic and reporting status of a Field Replaceable Unit  
5 (FRU);

Fig. 74 is a table giving an OCP Template;  
Fig. 75 is a table giving an OCP 8P Example;  
Fig. 76 is a table giving an OCP Button Label Example;  
Fig. 77 is an OCP Switches Block Diagram;  
10 Fig. 78 is a table giving Handling of Miscellaneous Commands;  
Fig. 79A, Fig. 79B, and Fig. 79C is a table giving CLI commands;  
Fig. 80 is a table giving settings of Modem Knobs;  
Fig. 81 is a table giving Modem Knobs settings for connection to a CLI port;  
Fig. 82 is a table giving Operation Limitations in a degraded system;  
15 Fig. 83 is a table giving Data Base Grouping;  
Fig. 84A is a first table giving Fields of a Partition Data Structure;  
Fig. 84B is a second table giving Fields of a Partition Data Structure;  
Fig. 85 is a block diagram of PBM Hardware;  
Fig. 86 is an Overview of PBM processes;  
20 Fig. 87 is a table giving Error Codes;  
Fig. 88 is a block diagram of a shared RAM communication;  
Fig. 89 is a block diagram of an MBM to CMM communication;  
Fig. 90 is a block diagram of an example of MBM forwarding;  
Fig. 91 is a block diagram of an example of CMM forwarding;  
25 Fig. 92 is a CMM COM port connection;  
Fig. 93 is a block diagram of a Telnet session;  
Fig. 94 is a field diagram of a Request message;  
Fig. 95 is a field diagram of a Response message;  
Fig. 96 is a field diagram of a Train message header format;  
30 Fig. 97 is a table giving a list of LAN formation group messages;

Fig. 98 is a table giving a list of messages in a reliable message group;  
Fig. 99 is a table giving a list of messages in a system discovery group of messages;  
Fig. 100 is a list of messages in a partition control group of messages;  
Fig. 101 is a table giving a list EV7 set up group of messages;  
5 Fig. 102 is a table giving a list of cable test group messages;  
Fig. 103 is a table giving a list of virtual console group messages;  
Fig. 104 is a table giving a list of firmware load and upgrade group of messages;  
Fig. 105 is a table giving a list of environmental retrievable group of messages;  
Fig. 106 is a table giving a list of FRU data group of messages;  
10 Fig. 107 is a table giving a list of error logging group of messages;  
Fig. 108 is a table giving a list of OS watchdog timer group of messages;  
Fig. 109 is a table giving a list of date/time group of messages;  
Fig. 110 is a table giving a list miscellaneous messages.

#### DETAILED DESCRIPTION OF AN ILLUSTRATIVE EMBODIMENT

15 Fig. 1 is a schematic block diagram of a symmetrical multiprocessor (SMP) system 100 comprising a plurality of processor modules 200 interconnected to form a two dimensional (2D)-torus mesh configuration. Each processor module 200 comprises two central processing units (CPUs) with connections for two input/output (I/O) ports along with 6 inter-processor (IP) network ports. The network ports are preferably referred to as North (N), South (S), East (E) and West (W) compass points. The North-South (NS) and East-West (EW) compass point connections create a (Manhattan) grid. Additionally, the outside ends of the mesh wrap-around and connect to each other. I/O traffic enters the 2D torus via I/O channel connections between the CPUs and I/O subsystem 150. Each compass point is coupled to an IP channel that comprises 32 bits of data and a 7-bit ECC code for a total of 39 bits of information transfer.

20 25 Fig. 2 is a schematic block diagram of the dual CPU (2P) module 200. As noted, the 2P module 200 comprises 2 CPUs with connections 210 for the IP (“compass”) network ports and an I/O port 220 associated with each CPU. The 2P module 200 also includes power regulators 230, system management logic 250 and memory subsystem 300 coupled to 2 memory ports of each CPU, wherein each CPU may have up to 16 gigabytes (GB) of memory per processor using

512 megabit (Mb) RDRAMs. In accordance with an aspect of the present invention, the system management logic 250 cooperates with a server management system to control functions of the SMP system. Each of the N, S, E and W compass points, along with the I/O and memory ports use clock-forwarding, i.e., forwarding clock signals with the data signals, to increase data transfer rates and reduce skew between the clock and data.

Each CPU is preferably an EV7 processor that includes an EV6 core, an I/O interface and 4 network ports. In the illustrative embodiment, the EV7 address space supports up to 256 processors and 256 IO7s in 16 GB mode. In 32 GB mode, the EV7 supports up to 128 processors with memory. The EV6 core preferably incorporates a traditional reduced instruction set computer (RISC) load/store architecture. In the illustrative embodiment described herein, the EV6 cores are generally the Alpha® 21264 processor chips manufactured by Compaq Computer Corporation®, with the addition of a 1.75 megabyte (MB) an internal cache and CBOX, the latter providing integrated cache controller functions to the EV7 processor. However, it will be apparent to those skilled in the art that other types of processor chips may be advantageously used.  
10 The EV7 processor also includes a RBOX that provides integrated routing/networking control functions with respect to the compass points. The EV7 further includes a ZBOX that provides integrated memory controller functions for controlling the memory subsystem.  
15  
20

The memory subsystem 300 is preferably implemented using RAMBUS technology and, accordingly, the memory space is generally divided between 2 RAMBUS controllers. However, an EV7 processor can operate with 0, 1 or 2 RAMBUS controllers.

Fig. 3 is a schematic diagram of the SMP memory subsystem 300 illustrating connections between the EV7 and RAMBUS memory modules (RIMMs 310). Software configures the memory controller logic (ZBOX 320) within the EV7 and the logic on each RIMM 310 before testing and initializing memory. Specifically, the memory subsystem components include 2 RAMBUS memory controllers (not shown) within the ZBOX 320, a RIMM 310 containing RDRAM memory devices, a serial I/O (SIO 330) channel to the RDRAMs of the RIMMs 310, serial presence detect (SPD) logic (EEPROM data) via an I<sup>2</sup>C bus 350, and a CPU management module (CMM) field programmable gate array (FPGA 360) that interfaces between a CMM (not shown) and the EV7 processor.  
25

Fig. 4 is a schematic block diagram showing the RAMBUS memory organization 400. Both EV7 memory controllers (ZBOX0/ZBOX1) and the RDRAMs 410 contain programmable elements which configure the addressing and timing of each RDRAM on each channel. The RIMMs are visible to system software via 3 separate paths: the SPD logic (via the I<sup>2</sup>C bus 350), the SIO channel 330 and the RAMBUS channels 380. The SPD path provides a data structure contained within a serial EEPROM located on each RIMM 310. The SIO path provides access to programmable timing and configuration registers contained on the RDRAM devices. The RAMBUS channel is the row/column serial data path to the RDRAMs. Each channel is connected to a physical RIMM that may contain up to 32 RDRAM devices.

Fig. 5 is a schematic block diagram of an IO7 device 500 that provides a fundamental building block for the SMP I/O subsystem. The IO7 is preferably implemented as an application specific integrated circuit (ASIC) using IBM SA27E ASIC technology and packaged in a 748-pin column grid array CCGA package. Each EV7 processor supports one I/O ASIC connection; however, there is no requirement that each processor have an I/O connection. In the illustrative embodiment, the I/O subsystem includes a PCI-x I/O expansion box with hot-swap PCI-x and AGP support. The PCI-x expansion box includes an IO7 plug-in card that spawns 4 I/O buses.

The IO7 500 comprises a North circuit region 510 that interfaces to the EV7 processor and a South circuit region 550 that includes a plurality of I/O ports 560 (P0-P3) that interface to standard I/O buses. An EV7 port 520 of the North region 510 couples to the EV7 processor via 2 unidirectional, clock forwarded links 530. Each link 530 has a 32-bit data path that operates at 400 Mbps for a total bandwidth of 1.6 GB in each direction.

In accordance with an aspect of the present invention, a cache coherent domain of the SMP system extends into the IO7 and, in particular, to I/O caches located within each I/O port of the IO7. Specifically, the cache coherent domain extends to a write cache (WC 562), a read cache (RC 564) and a translation buffer (TLB 566) located within each I/O port 560. As described further herein, the caches function as coherent buffers in that the information contained within these data structures are not maintained for long periods of time.

Referring again to the embodiment of Fig. 1, the 2D-torus configuration of the SMP system 100 comprises sixteen (16) EV7 processors interconnected within two 8P drawer enclosures

600. Specifically, there are four (4) 2P modules 200 interconnected by a backplane within each enclosure 600. This configuration is scalable by powers of 2 (EV7 processors) up to a total of 256 (or preferably 128) processors. In the illustrative embodiment, four (4) 8P drawers may be mounted within a standard 19-inch rack (2 meters in length) as shown in Fig. 6. Mounting 4 8P  
5 drawers in a single rack creates a substantial cabling problem when interconnecting the 32 processors within the 2D-torus configuration and when coupling the processors to the I/O subsystems via the IO7 devices 500 associated with the processors. In accordance with another aspect of the present invention, an efficient means for interconnecting cables among the 8P drawers of a fully-configured, 19-inch rack is provided.

10 Fig. 7 is a schematic block diagram of an I/O drawer 700 of the SMP system which includes a first I/O riser card 710 containing an IO7 500, a connector 720 coupling the IO7 to the EV7 processor and a plurality of I/O buses. The speed of the I/O buses contained within the I/O drawer is a function of the length and the number of loads of each I/O bus. The I/O drawer is divided into two parts: a hot-plug region 730 and an embedded region 750. In the illustrative embodiment, there is a dedicated slot 760 adjacent to the I/O riser card 710 within the embedded region 750 that is dedicated to a 4x AGP Pro graphics card. Also included within the embedded region 750 are 3 standard, 64-bit PCI card slots 772-776, two of which may be occupied by additional AGP Pro cards. Otherwise, these PCslots are available for embedded I/O card options. For example, an I/O standard module card 780 may be inserted within one of the PCI slots 772-  
15 776.  
20

Each I/O drawer 700 also includes power supplies, fans and storage/load devices (not shown). The I/O standard module card 780 contains an IDE controller for the storage/load devices, along with a SCSI controller for those devices and a universal serial bus that enables keyboard, mouse, CD and similar input/output functions. The embedded region 750 of the I/O  
25 drawer is typically pre-configured and not configured for hot-swap operations. In contrast, the hot-plug region 730 includes a plurality of slots adapted to support hot-swap. Specifically, there are 2 ports 732-734 of the hot plug region dedicated to I/O port one (P1 of Fig. 5) and 6 slots 738-748 dedicated to I/O port two (P2 of Fig. 5). Likewise, the dedicated AGP Pro slot 760 comprises port three (P3) and the 3 standard PCI slots 772-776 comprise port zero (P0). The I/O  
30 buses in the hot-plug region 730 are configured to support PCI and PCI-x standards operating at

33 MHz, 66 MHz, 100 MHz and/or 133 MHz. Not all slots are capable of supporting all of these operating speeds. In another aspect of the present invention, a technique is provided that enables all slots (under certain configurations) to support all operating frequencies described above.

Also included within the I/O drawer 700 and coupled adjacent to the IO7 is a PCI back-plane manager (PBM 702). The PBM 702 is an integral part of a platform management infrastructure as described further herein. The PBM is coupled to a local area network (e.g., 10 base 100 Ethernet) by way of another I/O riser board 790 within the I/O drawer. The local area network (LAN) provides an interconnect for the server management platform that includes, in addition to the PBM, a CMM located on each 2P CPU module and an MBM (Multiprocessor computer system Backplane Manager) located in each 8P drawer. In a preferred embodiment of the invention the Ethernet LAN comprises twisted pair Ethernet media coupling to an Ethernet hub. Note that the cable coupling the IO7 to the EV7 on a 2P module may be up to 6 meters in length.

Fig. 8 is a schematic block diagram of the server management platform 800 for the SMP system. The server management comprises a 3-tier management scheme. At the lowest level, each 2P module 200 has a plug-in, CPU management module (CMM 810) that provides power and initialization control for the local 2P module. The CMM also interfaces directly to both EV7 processors via serial links 820 and provides debug, initialization, error collection and communication support to a higher, intermediate level of the service management hierarchy.

The MBM 840 is preferably an independent plug-in card within an 8P drawer 600. Each CMM 810 on each 2P module 200 within an 8P drawer 600 communicates with an MBM 840 through a point-to-point serial connection 845 that is preferably implemented in etch so as to obviate the need for a cable connection. In the illustrative embodiment, each MBM controls 4 CMM devices within an 8P drawer.

The MBM 840 spawns a server manager network port that is connected to a service management LAN hub. The MBMs 840 preferably communicate with the PBMs 702 in the I/O drawers via a TCP/IP protocol over a server management LAN. In the illustrative embodiment, the server management platform is preferably implemented as a 10 base 100 (Ethernet) LAN, although similar types of local area network implementations, such as Token Ring or FDDI, may be advantageously used with the system.

A personal computer (PC) or similar network device connected to one of the ports of the service management LAN hub serves as a server management console (SMC 850). The SMC 850 provides the highest level of server management and, to the end, executes a platform management utility that provides a unified view of the entire SMP system for purposes of controlling the system, even if the system is divided into multiple hard partitions. From a physical implementation, the MBMs, PBMs and SMC are coupled to the service management hub; however, logically they are interconnected by the LAN.

The server management platform is used to bring up (“boot”) the SMP system and create partitions. As used herein, a hard partition is defined as hardware resources capable of supporting a separate instance of an operating system. In addition, the server management platform facilitates hot-swap (insert/delete) of hardware resources into/from the system. For example, assume it is desirable to dynamically remove a 2P module 200 from the SMP system 100. The SMC 850 instructs the appropriate MBM 840 which, in turn, instructs the appropriate CMM 810 on the 2P module to power down its regulators 230 in preparation of removal of the module. It should be noted that the SMP system may “boot” and operate without a functioning SMC, but reconfiguration and complete system visibility may be lost if redundant SCMs are not connected and the single SCM fails or is otherwise disconnected.

All console functions for the SMP system are provided by the SMC node 850 on the LAN. In the case of a hard partitioned system, a logical console is provided for each hard partition. In the illustrative embodiment, the SMP system may be expanded up to 256 processors (within, e.g., 32 8P drawers) and 256 I/O drawers, wherein each I/O drawer 700 includes a PBM 702 and each 8P drawer 600 includes an MBM 840. Therefore, the server management platform is preferably implemented as a local area network to accommodate such expandability. In an alternate embodiment, the SMC may be implemented as a web server and the server management LAN may function as a virtual private network (VPN). In accordance with this embodiment, system management of the SMP system may be controlled remotely over a network, such as the Internet, through a firewall at the SMC console station.

#### Virtual Channels

The SMP system comprises a plurality of virtual channels including a request channel, a response channel, and I/O read channel, and I/O write channel and an error channel. Ordering within a processor with respect to memory is achieved through the use of memory barrier (MB) instructions, whereas ordering in the I/O subsystem is done both implicitly and explicitly. In the case of memory, references are ordered at the home memory of the cache line data in a directory in flight (DIF) data structure (table) of the EV7.

In the I/O subsystem, write operations are maintained in order relative to write operations and read operations are maintained in order relative to read operations. Moreover, write operations are allowed to pass read operations and write acknowledgements are used to confirm that their corresponding write operations have reached a point of coherency in the system. Ordering in the I/O subsystem is important from the perspective of any two end points. For example, if processor (EV7a) communicates with its associated IO7 (IO7a), then all operations must be maintained in order. However, communication between another processor (EV7b) and IO7a is not maintained in order. If ordering is important, another mechanism, such as semaphores between processors, must be utilized.

### Deadlock Avoidance

Two types of deadlock may occur in the SMP system: intra-dimensional and inter-dimensional deadlock. Intra-dimensional deadlocks can arise because the network is a torus and the wrap-around path can cause a deadlock cycle. This problem is solved by the use of virtual channels. Inter-dimensional deadlocks can arise in any square portion of the mesh. These cycles can be eliminated if messages route all in one dimension before routing any in the next dimension or in dimension order. For example, if all messages traversed in the East-West (EW) direction before traversing in the North-South (NS) direction, no deadlock cycles can be generated because there is never a dependency from NS channels to the EW channels. The content of an RBOX configuration register selects whether the NS or EW channels are primary.

Dimension-order (i.e., deadlock-free) routing requires that a message route along a fixed path from source to destination. However, in some cases there may be multiple minimum-distance paths from source to destination. In this case, it is desired to select a path from source

to destination that encounters the least network connection. This is called “minimal adaptive routing”.

The EV7 processor allows for adaptive routing that is deadlock-free. Buffering is allocated for a deadlock-free network. Disposed over the deadlock-free network is an adaptive network of buffers. A message can travel on the adaptive buffers until it encounters a situation that might lead to deadlock. In this case, a message may exit the adaptive buffers and enter the deadlock-free network. Due to the nature of the deadlock-free network, that network can always make forward progress. Since the messages in the adaptive network may always drain into the deadlock-free network, the system should never deadlock. In effect, the adaptive network has dependencies on the deadlock-free network but not vice versa. A message in the adaptive buffers can re-enter the deadlock-free network at any point, just like any new message.

The SMP system allows adaptive routing to be performed based on the dynamic load in the network. However, the system is still deadlock-free because the deadlock-free network is always available. The majority of buffering is allocated to the adaptive network. There is minimal buffering in the deadlock-free network; that is, there is sufficient buffering to eliminate cyclic dependencies. The RBOX may (with the exception of I/O channel references in the normal case) start a message in either the adaptive or deadlock-free networks. When an adaptive message is blocked in the EV7 RBOX (router) due to lack of buffering, the message is converted to the deadlock-free network as space becomes available. It is also possible for the EV7 processor to convert from the deadlock-free network back into the adaptive network. Basically, the EV7 decides at each “hop” of the 2D-torus which buffer type it can use.

Generally, the following rules are followed on EV7 processors for a message to traverse the deadlock free network: (1) the message is routed in dimension-order from its current location, and (2) the message selects a virtual channel on each dimension-change. The EV7 can handle other special cases that violate these rules, such as an “L-shaped” system as well. It should be noted that messages in the I/O virtual channel are never routed adaptively.

The header of an EV7 message contains information indicating the direction that the message may take in each of the two dimensions and a value for each dimension. When the value for a given dimension equals “WHOAMI” (a stored value) of an EV7 processor at which

the message has arrived, the message is assumed to have traveled sufficiently far in that dimension. When the WHOAMIs of both dimensions equal both values contained in the message, the message has reached its destination. A routing table (RBOX\_ROUTE) in the RBOX holds the values that are sent along with each message, as well as the directions that a message should travel in each dimension. The routing table is consulted once as each message is sent, preferably at the source of the message. The information from the routing table is all that is required to find the destination processor. The routing table at each processor contains 276 entries, one entry for each processor in the system (plus one for each sharing mask bit).

#### I/O DMA Access and Exclusive Caching

The IO7 may perform direct memory access (DMA) accesses to the EV7 system memory by way of either exclusive caching or time-outs. A DMA device is contained within the IO7 and is configured to service I/O bus read and write operations. For a DMA write stream, a first way to prefetch data in multiple blocks is via a stream of read modify request (ReadModReq) commands. The second is via a stream of invalidate-to-dirty request (InvaltoDirtyReq) commands to gain exclusive access to the block (presumably to write the entire block). The InvaltoDirtyReq commands require that the write operations be full-block writes.

For a DMA read stream there are two ways to prefetch data in multiple blocks, depending on the ordering required by the DMA device. The most efficient way is to use a stream of fetch requests (i.e., non-cacheable fetch) commands, while another way is to use a ReadModReq command to obtain exclusive access to the block (often to write a portion of the block). The advantage of this latter way is that the I/O device can implement a sequentially consistent read stream since the exclusive access forces order. A disadvantage involves generation of Victim-Clean messages to release exclusive access to the block. Multiple DMA devices that attempt to access the same block at the same time must be serialized, as will a processor and a DMA device.

When using the DMA access and exclusive caching technique, the DMA device is expected to force the eviction of a data block (cache line) soon after receiving a forward for the cache block. The IO7 may exclusively cache copies of blocks for long periods of time. If a processor or another IO7 requests a copy of the block, the directory determines that the IO7 is

the exclusive owner of the block and forwards the request to the IO7. When this happens, the directory expects to eventually receive both a ForwardMiss and a Victim (or VictimClean) response.

When the IO7 uses exclusive caching to access DMA requests, it should respond with  
 5 ForwardMiss messages to every received forward request. The following is also required: (1) any currently cached blocks/TLB entries that could match the address in the forward message must be marked for eventual eviction (after a time-out); and (2) any currently pending miss addressed file (MAF) entries that could possibly match the address must be marked so that the block eventually gets evicted after it returns. It should be noted that the receipt of a forward  
 10 message does not imply that the IO7 currently holds a copy of the block. That is, a victim may be on its way from the IO7 to the directory before the IO7 receives the forward message. Note also that this scheme allows the IO7 to (exclusively) cache copies of scatter-gather maps or I/O TLB entries.

When using the time-out technique, the DMA device is expected to evict blocks soon after it obtains exclusive access to the block. This allows the IO7 to ignore the forward messages.  
 15 When the IO7 uses this mode to access DMA, it should respond with a ForwardMiss response to every receive forward request and otherwise ignore the forward message.

### I/O Space Ordering

The EV7 processor supports the same I/O space ordering rules as the EV6 processor: load  
 20 (LD)-LD ordering is maintained to the same IO7 or processor, store (ST)-ST ordering is maintained to the same IO7 or processor, LD-ST or ST-LD ordering is maintained to the same address, and LD-ST or ST-LD ordering is not maintained when the addresses are different. All of these ordering constraints are on a single processor basis to the same IO7 or processor. Multiple loads (to the same or different addresses) may be in flight without being responded to, though  
 25 their in-flight order is maintained to the destination by the core/CBOX and the router. Similarly, multiple stores (the same or different addresses) can be in flight.

The EV7 processor also supports peer-to-peer I/O. In order to avoid deadlock among peer IO7 “clients”, write operations are able to bypass prior read operations. This is required because read responses cannot be returned until prior write operations have completed in order to

maintain PCI ordering constraints. By allowing the write operations to bypass the read operations, it is guaranteed that the write operations will eventually drain, thereby guaranteeing that the read operations will eventually drain.

5

### Partitions

A domain is defined as a failure unit. A domain may constitute from one to many processors. The SMP system can be partitioned into domains via interprocessor register (IPR) settings. These domains provide varying degrees of isolation, survivability and sharing between 10 domains, such as hard partitions, semi-hard partitions, firm partitions and soft partitions.

In a hard partition, there is no communication between domains. In this type of system, an EV7 processor, memory or I/O failure does not affect another domain. Each domain can be individually reset and booted via separate consoles. Proper RBOX\_ROUTE, RBOX\_CFG and RBOX\_\*\_CFG settings are the primary requirement to establish a hard partition.

15 A firm partition allows domains to share a portion of its memory, the “global” memory, which is distributed throughout the domains. The local memory for each domain still resides within each domain. The EV7 processor can prevent domains from accessing local memory in other domains. An EV7 processor, memory or I/O hardware failure in one domain may cause corruption or fatal errors within the domain containing the failure. A hardware failure may also 20 cause corruption or failures in other domains. The proper settings in the RBOX\_ROUTE, RBOX\_\*\_CFG, CBOX Access Control, CBOX Local Processor Set and CBOX Global Processor Set IPRs are the primary requirement to set up a firm partition.

25 A semi-hard partition is a firm partition with some additional restrictions and hardware reliability assurances. It requires that all communication within a domain must stay within the domain. Only sharing traffic to the global memory region may cross domain boundaries. Hardware failures in one domain can cause corruption or fatal errors within the domain that contains the error. Hardware failures in any domain can also corrupt the global region of memory. However, hardware failures in one domain will not corrupt the local memory of any other domains

provided the local and global sets have been properly defined in the CBOX Local Processor Set, CBOX Global Processor Set, and CBOX Access Control IPRs and those IPRs are configured to deny access from remote processors. In addition, corruption should not occur provided the RBOX\_ROUTE configuration correctly directs local traffic with a local domain, inval sets do not cross domain boundaries and the time-out values are established to time out inter-domain channels before timing out intra-domain channels, as indicated in the time-out ordering.

A soft partition allows for all communication to cross domain boundaries. The domain is strictly a global software concept in this case. The partitions can share a global portion of the memory. Each domain has a region of local memory that the other domains cannot access. But a hardware failure in one domain may cause corruption in any other domain in a soft partition. The proper settings in the CBOX Local Processor Set, CBOX Global Processor Set and CBOX Access Control IPRs are the primary requirement to set up a soft partition.

## Server Management Subsystem

The Server Management subsystem forms an omniscient view of the system. This provides for logical partitioning, physical environment monitoring, power control, retrieval of saved error state, and console communications. The term “subsystem” is used to include both the hardware and firmware components that accomplish these functions.

Topics discussed include the following.

Server Management Subsystem Overview, describing both the hardware architecture and the general functionality of the Server Management Firmware.

Firmware Requirements, itemizing the required functions of the Server Management Firmware. These will be grouped by related area. Each area contains some amount of introductory text. Each requirement is contained in it's own numbered list item.

Firmware Top Level Architecture, describes the architectural decomposition of the listed requirements into the firmware components running on each hardware entity.

Reference documents describing the Alpha systems include *Alpha System Reference*

5 *Manual* (SRM), Sites, Witek, et. al., published by Compaq Computer Corporation, all disclosures of which are incorporated herein by reference.

### Server Management Subsystem Overview

10 The Server Management Subsystem is a distributed system of microprocessors which cooperatively control, monitor, and support the Multiprocessor computer system EV7 CPUs and I/O.

### Server Management Hardware Architecture

15 Turning now to Fig. 9, a block diagram 900 of Multiprocessor computer system Server Management Hardware is shown. The Server Management microprocessors communicate with each other via point to point connections and a private Local Area Network (LAN) 902. A Multiprocessor computer system is composed of dual EV7 CPU modules 200 interconnected via a backplane. Each EV7 CPU can optionally be connected to a PCI I/O drawer. This forms a modular, building block, oriented system. The Server Management hardware parallels this structure. Each dual EV7 CPU module contains a microprocessor called the CPU Module Manager 810 (CMM). Each backplane 904 of 4 dual processor modules contains a Multiprocessor computer system Backplane Manager (MBM) microprocessor. Each PCI I/O drawer 700 contains a PCI Backplane Manager 702 (PBM). The details of each of these hardware components 20  
25 are described below.

#### CPU Module Manager (CMM)

The CMM 810 communicates with the pair of EV7 processor chips, as well as performing serial RDRAM EEPROM I/O, thermal and power management functions. It is directly connected, via a serial communication path, to the local MBM. The CMM 810 can power the EV7 CPU module on and off, reset and halt the EV7 processors individually, and control EV7 Power

5 On Self-Test (POST) and initialization. This module includes a flash ROM containing CMM, FPGA, and EV7 code. The FPGA chip is a Field Programmable Gate Array chip and contains logic circuitry as described further below.

#### Multiprocessor computer system Backplane Manager (MBM)

10

The MBM communicates with each of the 4 CMMs in the local backplane via a point-to-point serial connection 902. It also performs thermal and power management functions at the backplane level. It communicates with its peer MBM modules and PBM modules via a twisted pair Ethernet LAN 902. The MBM includes a flash ROM for MBM and EV7 code as well as

15 non-volatile data storage.

#### PCI Backplane Manager (PBM)

20

The PBM 702 is responsible for the thermal and power monitoring and control of the PCI drawer 700. It communicates with the MBMs 840 in the system via a twisted pair Ethernet LAN 902. The PBM includes a flash ROM for PBM code and non-volatile data storage.

#### Server Management Firmware Overview

25

The Server Management Firmware is a distributed software system that helps orchestrate many facets of the hardware control of the system. It provides the user interface and platform level view, and thus a central point of monitoring and control. It also provides a network packet level interface that can be used by Platform Management applications.

30

Historical Perspective

Traditionally, AlphaServer systems have provided the “Console” interface, as described in the SRM, as the user interface prior to operating system boot. The SRM firmware runs on the Alpha CPU and is loaded from a flash ROM in processor I/O space. It provides commands a 5 Console Command Line Interface (CLI) prompt (>>>) to allow interrogation of the system configuration, test, reset. The Console also allows the setting of environment variables that alter the behavior of hardware and/or software. The Console is available via a serial line or a graphics video display. Additionally, some platforms have provided an optional “Remote Management 10 Controller”, or RMC. This RMC is generally tied into the platform's power & thermal monitoring and control circuitry and provides basic functions, such as “show” commands, reset, and halt. The RMC is connected between the operator terminal and the Console serial line. The operator communicates with the RMC via either an Operator Control Panel (OCP) or a CLI which is reached by typing a special “escape” on the system console serial line. There is generally little interaction between the SRM console firmware and the RMC, other than retrieval of thermal data 15 and a few non-volatile settings. Once the SRM Firmware is running, the RMC is in a passive state. Once system software is running, the SRM Firmware is also in a suspended state, until such time as the Operating System halts and returns control to the SRM Firmware.

A previous Server Management Subsystem, referred to as Alpha Server GS 320, provides 20 the next level of integration between the control and monitoring of the overall system and the execution of firmware / software on the Alpha CPU. Alpha Server GS 320 is a trademark of Compaq Computer Corporation for a computer system using Alpha processors. System software & SRM firmware can communicate with the Server Management microprocessors via a shared 25 RAM structure and a “master” System Control Module (SCM). The subsystem is an active participant in hot-plug and hard partitioning operations, as well as environmental monitoring and asset management. The SCM is connected in the path of the serial console terminal, as with the previous RMC implementations. In support of logical partitioning, the system provides for multiple SCMs, although only one is the master of the microprocessor interconnect at any given time. External hardware/software is required to provide a single point of control for the multiple 30 system console serial lines.

## Present Invention

The present invention uses Server Management firmware providing an integrated solution between the Alpha firmware and system software and the Server Management functions. This includes removing the restrictions surrounding serial communication lines, loading EV7 firmware from the Server Management subsystem rather than from a fixed I/O device, performing virtual console terminal connections, providing a global view of the system configuration, integrating partitioning and hot-plug operations, and catastrophic error state collection. The system provides a framework that can support these functions, and additional functionality that has not been implemented as of yet.

## External Interfaces

This section describes the hardware and software interfaces that are external to the Server Management Subsystem. The requirements section references these interfaces.

### Private LAN

The Server Management processors communicate over a private local area network (LAN) 902. It is separate from any network that the customer connects to the Multiprocessor computer system.

### Customer LAN

The customer LAN is a local area network that is used by the customer to communicate with their Multiprocessor computer system(s). Customer LANs are often connections to the worldwide Internet, and are made through Network Interface Cards supported by the PCI buses 732-738, etc. in the I/O drawers 700.

## Platform Sensors

Platform sensors include:

- Thermal sensors to read the temperature within an enclosure space
- 5 Voltage sensors which measure electrical voltages of power supply outputs
- Discrete sensors which indicate the good/bad state of fans & power supplies
- Discrete sensors which indicate the state of switches
- Discrete sensors which indicate the presence of components or modules

## 10 Platform Controls

Platform controls include:

- Power supply on/off
- Fan speed

## 15 Operator Control Panel (OCP)

There is one Operator Control Panel 922 (OCP) per 8P backplane. The OCP 922 is a  
20 module containing :

- an alphanumeric display
- discrete switches for halt and reset

## 25 Flash Image

There are flash ROMs associated with each microprocessor. These flash ROMs contain  
code & data storage for the following components:

- Microprocessor code
- FPGA programming data
- EV7 SROM

EV7 XSROM  
SRM Firmware & PAL  
Non-volatile Configuration data  
Catastrophic error log data

5

### SRM Firmware

The SRM firmware is the EV7 code that runs the “Console”. It is dependent upon the Server Management Firmware for:

10        Loading  
            Retrieving some configuration information  
            Retrieving current logical partitioning state  
            Performing logical partitioning allocation and deallocation  
            Console Virtual Terminal I/O  
15        Platform Sensor data

### PALcode

20        The PALcode component performs the functions required by the Alpha SRM. It makes use of the Server Management firmware for logging catastrophic error information.

### Serial Console

25        The serial console is a serial UART port, used for system debug of systems of greater than 2 processors, or as the system console on a dual processor Multiprocessor computer system.

### RIMM I<sup>2</sup>C Interface

The memory RIMMs include serial EEPROM components that contain the Serial Presence Detect (SPD) data required in order to configure the RAMbus memory subsystem. The EEROMs are connected via an I<sup>2</sup>C bus to the CMM.

5 EV7 CPUs

The Alpha EV7s 860 interface directly to the CMMs 810. This connection includes:

- Halt and Reset
- SROM clock and data
- 10 GPORT I/O bus
- BIST signal(s)

System Software

15 System software interfaces with the Server Management Firmware to:

- Obtain sensor data
- Allow agent software to issue commands to the Server Management Subsystem
- Obtain and save the current date/time via the SRM defined BB\_WATCH, as described herein below.

20 Operator Station

The operator station 922 is defined as the location running the Platform Management Utility.

25 This utility uses the Private LAN 902 to perform the Server Management operator functions.

Requirements

The description that follows is for the purpose of showing exemplary functional requirements of the Server Management Firmware. The Server Management Firmware will be referred to as “the firmware”. Any other firmware will be called out specifically. This specification describes the requirements as follows.

5

Initialization

Network Processing

Environment and Configuration

OCP and Console Traffic

10

EV7 Control

Error Logging

SRM Firmware Interface

Server Management Command Execution

Partition Management

15

Real Time Clock (BB\_WATCH)

Firmware Update

Platform Debug Utility

Initialization

20

The Server Management microprocessors are powered by an auxiliary power supply (Volts Auxiliary or VAUX) power supply, separate from the main system supply. This supply is present even in the absence of system power.

25

Upon application of VAUX:

1. Module self-test is performed on all Server Management processors (CMM, MBM, and PBM). Self-test status is stored by each Server Management processor for retrieval upon command.

2. There is a discovery process by which all members learn of the total population. This is accomplished via communication over the Server Management LAN.
3. The system supports the re-initialization, addition, or deletion of Server Management processors without affecting running operating system instances.
5. Upon completion of the Server Management Subsystem initialization, CPU and I/O initialization is begun, if the system master power switch is enabled / on.
5. The functions of system initialization and operating system bootstrap are independent of the presence of an Operator Station or Platform Management Utility.

10

## Network Processing

This section describes the requirements related to the distributed processing and communication on the private Server Management LAN 902. Each microprocessor is considered to be a "member" of the network.

1. A message packet protocol is used for communication between any PBM, MBM, CMM, and EV7 on the private LAN.
2. The addressing on the private LAN is constructed such that each member is addressed in a deterministic method, based upon its geographic location, physical node ID, or other software visible feature.
3. The MBMs and PBMs directly connected to the Ethernet LAN operate as peers. Traffic from subordinate CMMs is routed via its directly connected MBM.
4. The network protocol supports periodic announcement messages, which are used by each of the individual members to build a complete list of members.
25. Traffic on the private LAN is not be visible to the customer's LAN.

6. An optional Operator Station 930 may be connected to the private LAN, function as a member of the network, and communicate using a Platform Management Utility.
7. Address selection for multiple operator stations is automatic and transparent to the user, if desired.

5      Remote LAN Communication

Remote LAN communication through connection 932 is defined as any Server Management LAN traffic that originates from the Customer LAN (as opposed to a directly connected operator station).

10

8. Connection to a customer LAN through connection 932 is via a two port gateway. The gateway contains a single port on the private LAN and a single port on the Customer LAN.

9. The gateway is assigned a single IP address on the Customer LAN.

10. The gateway functions is a member of the Private LAN. A private LAN IP address is assigned via a deterministic method.

11. Only external commands (defined in a subsequent section) are accepted from the Customer LAN side). Illegal commands are rejected.

12. Commands are passed through the gateway to the Private LAN. Responses are returned from the Private LAN to the Customer LAN.

20 13. The gateway function implements a means of access control (e.g. username / password) from the Customer LAN side.

These features may be satisfied by a gateway function running on the Operator Station platform, or may be implemented in another embedded micro processor.

25      Environment and Configuration

This group of requirements involves the monitoring of the physical configuration, thermal, and electrical environment of the Multiprocessor computer system.

14. Power on/off sequencing for the entire system, as well as hot-plug capable modules & drawers, are coordinated by the firmware.

15. A list of all of the following hardware components, with associated asset information (revision, serial #) shall be maintained:

EV7 CPUs

RDRAM RIMMs

IO7s

CMMs, MBMs, PBMs

10 I/O Drawers

PSUs and Fans

16. The CMMs retrieve the RDRAM SPD data from each RIMM on all CPUs. Further, this data is provided to the EV7 XSROM code.

15 17. The BIST status of all EV7 CPUs in the system is stored.

18. The configuration data and/or sensor data is provided upon command.

19. The Primary EV7 CPU in each partition is interrupted upon configuration changes that add or remove hardware resources (CPU or I/O chassis).

20. Configurable parameters are maintained in non-volatile storage. In a system that contains multiple MBMs, there are multiple copies of this data.

21. A mechanism for the SRM firmware to store and retrieve non-volatile parameters (referred to as environment variables) is provided.

22. In the event that the value of a critical platform sensor enters a region defined as hazardous, steps are taken to power down the components within the domain of that sensor. Notification is made to the OCP and to any attached Platform Management Utility that such an event has occurred. The occurrence is logged.

23. A set of functions is provided which allow a discrete signal to be set, cleared, and checked from each side of an EV7 Interprocessor cable.

## Configurable Parameters

Operator settable parameters that are maintained by the firmware include the count of partitions, the LP-count and the identity of the EV7 Processors and I/O drawers belonging to each partition.

### OCP and Console Traffic

The OCP provides the lowest level control of the system and diagnostic feedback to the user. The traditional text based interface to an AlphaServer and any associated Remote Management Controller is considered “Console Traffic”. This group of requirements covers the OCP interface and ASCII character I/O between the user and:

- 15      The SRM Console in each logical partition
- System Software Console I/O via the SRM console port
- The 2P Serial Console / Debug Console

### OCP Interface

- 20      24. The alphanumeric display on each 8P backplane is used to display critical error and status messages.
- 25      25. The halt switch is programmable to be inoperative, halt all partitions in that 8P backplane, or halt all processors in all partitions.
- 26. The reset switch is programmable to be inoperative, reset all partitions in that 8P backplane, or  
25      reset all processors in all partitions.

### SRM Console Firmware Traffic

27. A mechanism is provided by which the SRM Console firmware, running in a given logical partition, performs ASCII character I/O to the Platform Management Utility. A connection to a given logical partition is defined as a session.

28. The Platform Management Utility provides a means to connect to the operator display and keyboard to a session running on each partition.

29. Characters of output are buffered in the absence of an established session.

30. SRM Console firmware messages leading up to the first user input prompt are stored as an audit trail for later retrieval by the Platform Management Utility. This audit trail is re-written each time that the logical partition is initialized by the Server Management firmware.

1031. Simultaneous character I/O sessions are supported.

#### System Software Console I/O via the SRM console port

15        Per the Alpha SRM, “Alpha console provides system software with a consistent interface to the console terminal, regardless of the physical realization of that terminal.”. The Multiprocessor computer system takes advantage of this to virtualize console character I/O. The SRM Console firmware utilizes the same interface as described in the section above to perform this function.

20

The features necessary for the SRM Console firmware to implement the GETC, PUTS, PROCESS\_KEYCODE, RESET\_TERM, SET\_TERM\_INT, and SET\_TERM\_CTL callbacks are described in the Alpha SRM.

25        2P Serial Console / Debug Port

A serial port connection, via the MBM is to be utilized for two purposes. In the case of a minimal configuration Multiprocessor computer system (2P, single partition) there may not be an

Operator Station. Also, for the purpose of system lab and/or software debug, it is desirable to have a hard-wired serial line connection closer to the EV7 CPU.

32. There are, in debug mode or for a 2 Processor system, a means provided to connect the SRM Console firmware traffic to a Serial Console line on a given MBM.

5 33. Given the same configuration requirements, this mechanism is supported as the System Software console I/O port.

## EV7 Control

10 The Server Management Subsystem orchestrates, via hardware and firmware, the power on sequencing of the EV7 CPUs.

34. The EV7 SROM code is loaded from flash into the CMM FPGA.

35. The EV7 XSROM code is loaded from flash via the CMM FPGA.

15 36. EV7 startup sequence with the SROM / XSROM is sequenced in a synchronized fashion, following the power-up flow.

37. A means to execute individual SROM / XSROM tests is provided.

38. The SRM Console code is loaded from the MBM flash to the EV7.

39. The EV7 CPU RESET sequence is performed upon command.

20 40. The ability to set and clear the EV7 HALT signal is provided via a command.

41. The proper power-on / off and reset sequence to initialize the EV7 is performed via the CMM.

42. A means of loading EV7 firmware (XSROM or SRM Firmware), as an alternative to the contents of the flash image, is provided for diagnostic purposes.

25

## Error Logging

The Server Management firmware provides a mechanism by which key information from catastrophic failure conditions can be collected for later analysis. It also provides functions to support storing specified FRU error data in that FRU's nonvolatile storage.

5

43. A means is provided for EV7 PALcode to signal errors designated as catastrophic (e.g. double error halts). Further, it shall provide for the non-volatile storage of such errors. This storage shall be cleared upon user command.

44. SROM / XSROM progress status and/or messages is stored as an audit trail, per EV7, for retrieval by the Platform Management Utility. This audit trail is re-written each time that the EV7 is initialized by the Server Management firmware.

10 45. A mechanism is provided to allow SROM, XSROM, or other diagnostic firmware to store failure date within FRUs which support it.

#### SRM Firmware Interface

15 This section describes the interdependencies between the SRM Console Firmware and the Server Management Subsystem. This interface provides the SRM Console Firmware with the ability to issue Server Management commands and access to the system state and configuration data. Console character I/O has been covered in a previous section.

20 46. The SRM Console Firmware executing in each logical partition instance has the ability to send and receive network packets on the Server Management LAN.

47. The SRM Console Firmware interface allows any valid Server Management command (described in another section) to be executed.

25 **Server Management Command Execution**

This is the set of requirements surrounding Server Management command interpretation and execution. There are two types of commands present on the Server Management LAN. Internal commands are those that provide communication between the microprocessors and/or the Plat-

form Management Utility, e.g. “print a character”. External are commands which transfer information or control to/from the outside world, e.g. “show FRU”.

48. The Platform Management Utility allows external Server Management commands to be  
5 issued by an operator.

49. Internal commands are provided for intercommunication between private LAN members.

#### External Server Management Commands

The table of Fig. 11 contains the set of required External Server Management Com-  
10 mands.

#### Internal Server Management Commands

The table of Fig. 12 contains the set of required Internal Server Management Commands.

#### Partition Management

20 The Server Management Firmware has a specific role to play in the area of Logical Partitioning. The Galaxy Configuration Tree Specification describes the role of the Server Management Subsystem in a partitioned system as well as the callbacks provided by the SRM Console Firmware.

25 50. The count and identities of hard partitions in the system is maintained.

51. The state of the current partition ownership of the hard partitionable entities (an “entity”) in the system is maintained. In an exemplary embodiment of the invention, the only hard partitionable entities in a Multiprocessor computer system are the EV7 CPUs. Placing a CPU in a hard partition also implies that its memory and associated IO7 are within that partition.

52. A mechanism to store the current partition count and ownership information in non-volatile storage is provided.

53. Partition initialization on power-up according to the partition information stored in non-volatile storage is performed.

5 54. A mechanism by which to provide the current partitioned state to the SRM Console Firmware running in any partition is provided.

55. A mechanism to allow the SRM Console Firmware, running in any partition, to change the ownership state of an entity is provided.

10 56. A mechanism by which an Operator can define which partitions may be allowed to allocate specific, free, entities is provided.

57. The primary CPUs of all partitions are notified of changes to the partition and ownership information, via interrupt. There is a mechanism for the SRM Console Firmware to enable/disable this interrupt upon request by System Software.

15 58. The logical and physical EV7 processor IDs are passed to firmware executing on the EV7 processor for use during configuration.

### Real Time Clock (BB\_WATCH)

20 The Alpha SRM defines the entity which stores the battery backed up date and time information as BB\_WATCH. System software uses the information stored in the BB\_WATCH to provide a consistent date/time across system reboots and power outages. This section describes the requirements for having Server Management firmware provide the BB\_WATCH functionality.

25 59. System Software running on a primary EV7 is provided access to read and write the BB\_WATCH data described in the Table of Fig. 13. The access is via the EV7 GIO port.

60. The Server Management firmware maintains a battery backed up timebase.

61. A unique set of BB\_WATCH data is provided for each partition.
62. The time for System Software to read or write the BB\_WATCH data is less than one second.

### Firmware Update

5

The contents of the Flash Image are distributed throughout the Server Management Subsystem. This section describes the requirements related to updating the various flash images.

63. A Fail Safe Loader (FSL), which executes in the absence of a valid firmware load, is implemented on each microprocessor. The Fail Safe Loader initiates a firmware update of the affected flash image. The FSL is protected by hardware to prevent unintended erasure.
64. The update of multiple flash ROMs is coordinated by the Platform Management Utility.
65. A discrete input, such as a jumper, is employed to prevent against accidental firmware update.

15

### Firmware Top Level Architecture

This section describes the top-level firmware architecture and allocation of the requirements from the previous sections.

- 20 Fig. 10 gives the Top Level Architecture of Firmware

The Server Management Firmware implementation contains the following components:

CMM Firmware 10,002

25 MBM Firmware 10,004

PBM Firmware 10,006

CMM Firmware

The CMM is responsible for the direct execution of all EV7 related tasks. In addition, it must satisfy the process Environment and Configuration requirements that relate to the Multi-processor computer system Dual CPU board. It directly executes the Server Management commands described in section 47 that are applicable to its operation.

### MBM Firmware

10

The MBM firmware is responsible for orchestrating the overall system initialization. As peers on the private LAN, each MBM participates in the network operations. By communicating with the PBMs and subordinate CMMs, it performs tasks required by the processor.

15

In cooperation with the other MBMs, the configurable parameters are maintained. The MBMs perform a routing operation in the routing of Console Traffic. The MBMs perform the partition management functions. It also directly executes the Server Management commands that are applicable to its operation.

### PBM Firmware

20

The PBM firmware participates with other PBMs and MBMs on the private LAN to satisfy the PBM requirements. It provides environment sensor data. It also directly executes the Server Management commands that are applicable to its operation.

25 Trademarks

VxWorks is a trademark of Wind River

AMD, is a trademark of Advanced Micro Devices.

Am186, is a registered trademark of Advanced Micro Devices.

Microsoft Word 97 and Microsoft PowerPoint 97 are registered trademarks of Microsoft Corporation.

LM80, is a trademark of National Semiconductor Corporation.

I<sup>2</sup>C, is a trademark of Philips Semiconductors.

## 5 Definition of Terms

Adaptive routing: the collection of paths advancing node-to-node in the same primary and secondary directions as the dimension-order routing. At each intermediate node it must be possible to advance in either direction until the dimension coordinate in a direction matches that of the destination.

10 Dimension-order routing: the shortest path connecting two nodes which proceeds first along the primary dimension and then along the secondary.

DRDRAM : DRAM memory chips that conform to the Direct RAMbus specifications. This new memory chip architecture provides extremely fast access times, 60nS to a closed page and 30nS to an open page. In addition to fast access time Direct RAMbus DRAMs support high bandwidth. A single RDRAM can support 1.6GB/sec of memory bandwidth.

15 Dynamic Duo : Multiprocessor computer system Dual processor CPU module. Includes two EV7 chips, their memory RIMMs, the 48V to DC converters, CPU server management logic, and connectors for six IP ports and two IOP ports.

EV7 : New version of the Alpha chip design based on EV68. The EV7 chip uses the EV68 core at its center and adds an on-chip memory controller, on-chip L2 cache (no module level cache support in this version of the Alpha chip), on-chip processor-to-processor and processor-to-IO router. The initial implementation uses 0.18um bulk CMOS with copper interconnect.

EV7x : EV7 chip implemented in a future CMOS process. 'x' denotes the process generation, i.e. EV78 is the EV7 chip produced with the next generation CMOS process.

25 EV8 : Next generation Alpha chip following EV7.

Group: A set of interconnected MBMs and PBMs.

Hard-Partition: A subset of the system's resources that is electronically isolated.

Initial hop: a routing option which allows a hop from the source node in any direction to an adjacent node. This option allows some connection of nodes in imperfect meshes.

IO7 : IO ASIC that attached directly to the EV7 IO port. The IO7 spawns four independent IO buses. Three of the buses are PCI-X and one of the buses is a 4x-AGP. This combination of buses provides for a very flexible IO subsystem.

IP : Inter-processor port on an EV7 chip. The IP ports are designated as North, East, South, and West relating to geographical location of the port in the interconnect mesh connecting processors together.

IOP : IO port on the EV7. This is the port that the IO7 ASIC connects to. The IOP has two unidirectional clock forwarded paths for communicating with the IO7 ASIC.

Members: MBMs and PBMs that belong to a group.

10 Partition: A subset of resources.

Partition Coordinator: A Server Management firmware task that manages a hard partition and all of its subpartitions.

Partition Database: A database of configuration information that is replicated among all the MBMs and PBMs.

15 PCI-X : New generation registered version of PCI bus. PCI-X runs at 133MHz, 100MHz, or 66MHz. PCI-X is backward compatible with PCI and will run at 33MHz or 66MHz when such a PCI device is plugged into the bus. The bus supports 32 bit and 64 bit devices.

Primary dimension: one of EAST-WEST or NORTH-SOUTH. This choice is the same for all EV7s in a hard partition.

20 RIMM : Memory DIMM utilizing DRAMs conforming to the RAMbus Inc.'s DRDRAM specification and RAMbus Inc.'s RIMM specification. RAMbus Inc. owns the specifications and is working with RAM vendors and memory suppliers to make RIMMs a commodity memory.

Secondary dimension: the other way (see primary dimension).

SIO : A new IO specification still being defined.

25 SMP : Symmetric multiprocessing. This is a tightly coupled shared memory multiprocessor system.

SRC routing: another deadlock-free routing method in which travel proceeds first along the secondary dimension.

Striping: Interleaving of the NUMA memories

**Sub-Partition:** Partitions within a hard partition that are not completely electronically isolated.

**Zone:** The subsets of hardware resources that results from a hub or link failure.

## Server Management

5        The Multiprocessor computer system Server Manager Firmware subsystem has two basic obligations : (i) to serve the user requests to act on Multiprocessor computer system hardware, and (ii) to automatically manage the Multiprocessor computer system hardware according to rules and constraints.

10      A user walks up to a laptop running the Partition Management Utility (PMU). He/she issues the command to power on a partition. The PMU sends the power on command to the PMU Server running on an MBM. The PMU Server forwards that command to the Partition Coordinator, which runs on an MBM. The Partition Coordinator sends a sequence of commands to other MBMs, CMMs, and EV7s to power up the hardware in the partition, to run the diagnostics, 15     to configure the routers, and to bring up the SRM console and operating system. The user's request is complete.

### Multiprocessor computer system Server Management Firmware Hardware Overview

Figure 15 shows hardware components of the Multiprocessor computer system Server Management subsystem, along with the associated firmware. The Multiprocessor computer system Server Management Firmware is the set of distributed hardware and software that facilitates managing a Multiprocessor computer system. The Multiprocessor computer system hardware is managed by a collection of microprocessors acting cooperatively. These micros are known as the CMM, the MBM, and the PBM. The CMM resides on the CPU module 810 and is affiliated with two EV7s. The MBM resides on the 8P backplane on MBM unit 840 and is a parent to the CMMs. The PBM resides in the IO box in PBM unit 702 and is a peer to the MBM.

The communication architecture is a 100baseT private LAN 902 and serial UART connections between the MBM and CMM Processors. The MBMs and PBMs have a direct connection to this LAN. The CMMs are connected to the MBM with a dedicated serial connection. The MBM routes the network messages for the CMM.

## 5 MBM Hardware Overview

Below is a brief overview of MBM hardware(16,000), as shown in Figure 16..

ELAN SC520 – compute engine for the MBM, a microprocessor 16,002

Quad Uart 16,004— each uart connects to a CMM via PPP

Flash 16,006— Stores firmware images for MBM, CMM, SRM, and data such as  
10 the partition information and data bases, the SRM environment variables, error logs

DRAM 16,008— Memory used to run the MBM firmware; ECC checked

NVRAM 16,010— a combination of flash and either CMOS or some larger  
15 NVRAM device. Frequently changing state and data is stored here.

I2C Interface 16,012— EEROM, fans, voltages, temperature

I2C Bus 16,014

TOY clock 16,016— Most importantly used as the clock for the EV7s controlled  
by this MBM

OCP display and switches 16,020— Displays information about the state of the  
20 CPU box. Switches are software managed

Rack and box thumbwheels 16,022, 16,024- the MBM derives a unique address  
from these; they also aid the human in identifying components

IP & IO Cable Connectivity 16,026— 14 cables are tested for expected connectivity and indicators help the human connect up the cables

Ethernet controller 16,028— All MBMs and PBMs are connected via private Ethernet 902. Management utilities are also share this LAN

Local terminal uart – in the lab, and for the 2P system, a simple connection is available not shown, in Fig. 16.

5 Debug uart – only used in the lab by developers, not shown in Fig. 16.

### Addressing the hardware

All MBMs are identical except for a uniquely identifying thumbwheel. The MBM uses this unique identifier to determine a unique IP address for itself and for 4 CMMs and 8 EVs. The IP addresses are reserved for the CMMs and EVs even if these components are not present.

10 Box and Rack Thumbwheels

Each 8P drawer and IO box as shown in Fig. 17 has a user settable thumbwheel with a 4 bit value (0-15). The thumbwheel value can be read by the PBM and the MBM. In addition, there is a thumbwheel identifying the rack in which the box resides with an 8 bit value (0-255). MBM thumbwheels 16,022, 16,024 indicate values of 0-3 on the 8P drawer and 0-7 on the rack.

15 PBM thumbwheels can have any value for the IO box and rack. These thumbwheels are used by the MBM and the PBM to derive a unique IP address on the SM 16,002, etc. The thumbwheel settings are available to the micros when the system is powered off, and always visible to the user.

20 Identification by thumbwheel is just one of the four naming schemes being used in an exemplary Server Management Subsystem implementation.

In addition, the rack and box identification is used to verify that InterProcessor cabling is correct.

### MBM Firmware Overview

The main duties of the Server Management firmware are:

25 Configure the partitions and associate the EV7s, IO7s, and memory to a partition

Provide each partition with individualized non-volatile storage and time services

Provide services to the CMM (Network proxy, time service, firmware repository...)

Power on and Power off components

Orchestrate diagnostics and perform cable testing

5 Offer console connections to each partition (virtual console)

The PBM duties are a subset of the MBM. A section dedicated to the PBM is given herein below.

10

### Operational Overview

#### MBM self management

The MBMs are the core of the Multiprocessor computer system Server Management Sub-system. They are parents to the CMMs and they service the Platform Management Utility. Every MBM needs to manage its network connections, its environmental, and serve its CMMs. Task that always run on each and every MBM:

Group protocol

Time Server

20 CMM TFTP Server

TFTP Client

Environmental Monitoring

## MBM LAN management

MBMs all have the same capability. But at any moment, some MBMs are responsible for more of the SM capabilities than others. As the configuration changes, the responsibilities of the MBMs change. This design allows for a robust architecture in the presence of faults. Since all MBMs are capable of performing any of the Server Management duties, an SM service can failover to any other MBM (or PBM). Additionally, some load balancing is achieved by distributing the labor. After each configuration event, the SM subsystem re-evaluates the task assignments, and restarts task as necessary.

Tasks that are instantiated on selected MBMs are shown in the Table of Fig. 18. In a degraded system that has broken into zones, there are 1 of these tasks for each zone: "Group Leader" and "PMU Server".

## Server Management Startup – an overview

When a Multiprocessor computer system is powered up, the MBMs and PBMs connect to the private LAN 902 and discover each other. They all participate in the "Group Membership Protocol". They agree on a "Group Leader" MBM/PBM to provide a consistent view of the replicated databases and to monitor the LAN membership. Since each MBM/PBM has a copy of the partition database, they each send their copy to the Group Leader for "reconciliation". The Group Leader distributes the authoritative copy of the database back to each member.

The Group Leader starts up the "PMU Server" task to handle (i) requests from the Platform Management Utilities (PMUs), (ii) error alerting, (iii) cable testing requests from partitions. Next, a "Partition Coordinator" task starts for each partition defined in the partition database. The Partition Coordinator task is the focal point for managing a partition (a partition is a component subset of the entire Multiprocessor computer system). The Partition Coordinator controls the MBMs/PBMs in the partition. It synchronizes diagnostic testing, and initiates the partition startup.

With all server management tasks ready, the DHCP server now starts up and offers network addresses to operator stations running PMUs. The Platform Management Utility(s) run on a separate system that requests a DHCP address and begins communicating with the PMU Server.

Each of these tasks is discussed in more detail in the following sections.

## Parts of the MBM

Below is a brief description of the items shown in Fig. 19.

5 CMM Network stack 19,002– UART driver 19,010, UDP19,006, PPP 19,008,  
TFTP 19,012. The MBM talks to the CMM over the serial lines using standard  
network protocols for much of its work.

SM Protocol – The Server Manager Protocol 19,014 is described. This is a private  
protocol.

10 TFTP 19,020 Server – Requests from the CMM for firmware that is stored in  
MBM flash are serviced by the parent MBM. Requests from an MBM for  
firmware is requested via BOOTP.

15 TELNET COM Forwarding – The console terminal interface is via a Telnet ses-  
sion which runs on the MBM. Characters are passed on further to the EV7 via  
the FPGA uart emulation hardware. Up to 16 TELNET sessions are possible on  
a single MBM (8 partitions x 2 COM ports per partition).

Flash Driver 19,022– The flash driver is accessed for reads and writes through a  
device driver. There is one flash driver running per MBM.

20 Time Services – The MBM provides the access to the TOY 19,024 for VMS and  
UNIX operating systems running on the EV7s. It also implements the synchro-  
nization of time across the partition. Each MBM runs the Time Server 19,026.

25 2P CLI & Debug CLI 19,028– In the 2P configuration, no laptop is required for  
management. A simple command interface is provided. This interface mecha-  
nism is also used as a debug mechanism on large configurations. Each MBM  
runs the Debug CLI. In the 2P system, there is only 1 MBM.

Modem 19,030– Dial out is for alert conditions. This is only run from one MBM, designated by a Server Management Knob. The connection to the modem is not a fault tolerant service since in an exemplary embodiment, the modem only has one connection into the system.

5 Cable Testing 19,032– User’s are assisted in connecting cables with LEDs and commands. Cables are also tested as part of diagnostics. Cable Testing is orchestrated by the PMU Server on request from either the PMU or a Partition Coordinator. All MBMs and PBMs participate.

10 OCP Manager 19,034– A small amount of information can be displayed on the OCP to indicate the configuration and errors. Each MBM has an OCP to utilize.

I2C Driver 19,036– The 8584 I2C controller is accessed for reads and writes through a device driver. It is used to access sensors, EEPROMs, the OCP, and sundries.

15 Environmental Monitoring 19,038 – Fans, temperatures, and voltages are monitored and notification of out-of-range conditions alerts the PMU.

Error Handling 19,040– Errors that the MBM detects (ECC, environment, power cycling,...) and other assorted errors are stored in the log.

Error Logging 19,042– a history buffer of errors is maintained. It can be retrieved by the PMU.

20 FRU callout 19,044– MBMs evaluate the results of diagnostics and alert the PMU about a failing component.

POST 19,046- Power On Self Test. The MBM does some self checking diagnostics. The POST runs on every MBM and PBM.

25 FSL 19,048– The Fail Safe Loader tries to recover from a corruption of the standard MBM flash image. An FSL exists on each MBM and PBM.

MBM Kernel – VxWorks, multitasking real image is used as the kernel.

Partition Coordinator & Ctable – A process runs for each partition that manages and synchronizes all the elements in the partition. There is one Partition Coordinator for each partition. An MBM may have from 0 to 8 Partition Coordinators running on it.

Group Protocol Server & Leader – The group protocol, runs on each MBM and PBM to ensure a coherent assignment of the Multiprocessor computer system resources. One MBM in the system is elected as leader to orchestrate this operation. Each MBM participates in the Group Protocol, but only one MBM is a leader per group. There may be more than one group in cases when a hardware failure divides the system.

PMU Server – One of the MBMs is elected to service the requests of the PMU using a well-defined IP address. In a split system there is a PMU Server for each “zone”.

TFTP Client & Server 19,050– Update images and test images are passed back and forth via TFTP. When an MBM’s image is corrupt, another MBM services a bootp request and provides the ailing MBM with a good image.

DHCP Server 19,052– One of the MBMs in the system provides IP address to PMUs joining the LAN.

LAN Network Stack 19,054– Ethernet driver, UDP, TCP, ARP, ICMP, DHCP, TFTP, Telnet, PING and the SM Protocol

#### Group Leader

The information used to configure a Multiprocessor computer system is stored in the flash on the MBMs and PBMs. There is no master MBM or PBM, but rather the configuration data is replicated on each MBM and PBM. This facilitates dynamic configuration changes and robustness during failures. And although there is no master, the LAN micros elect a leader, as shown in Fig. 20, to coordinate the system’s operation.

The MBMs and PBMs in the system coordinate by forming a set of micros, the group, that has a common view of the micros that exist. In the absence of failure, all the MBMs and PBMs are “members” of the group. The micro with the lowest thumbwheel identification is 5 chosen as the “leader”.

Each MBM and PBM has a copy of the partition configuration data. It is possible, due to a history of configuration changes, the replicated partition configuration data (sometimes referred to as the partition database) may differ from micro to micro. It is the job of the leader to 10 reconcile these differences and to distribute the reconciled database to all the members. The leader also starts the different management services, like the DHCP server and the PMU server.

A group forms or reforms whenever an MBM/PBM connects to/disconnects from the SM LAN (including failures, MBM/PBM crashes, etc.). At the time a group reforms, all existing 15 services are shut down, the leader is temporarily non-existent, and access to the database is locked. Each time the group reforms, the leader is again redetermined. If the lowest id MBM/PBM leaves the group, the next lowest valued MBM/PBM becomes the leader. If a new lowest id MBM/PBM connects to the LAN, it becomes the new leader. The new leader reconciles the database and also restarts the different management services, like the DHCP server and 20 the PMU server.

The service shutdown and restart behavior at group formation time is consistent throughout the relocatable services in the system. The Partition Coordinators tasks are relocatable (see section 0, “Partition Coordinator”). TELNET services are not relocatable because Telnet uses 25 TCP, a connection oriented protocol.

As shown in the table of Fig. 21, the power up flow is shown with emphasis on the role of the group. This flow can best be appreciated when we imagine an operator walking from one rack to the next and switching the breakers on each power supply to ON. MBMs/PBMs start 30 serially, just moments apart. The MBMs/PBMs form groups, reform groups, and run system di-

agnostics until the system reaches the phase where the SRM is running on a partition. After that point, changes in group membership (MBMs/PBMs coming and going) are treated as hot adds.

### The Platform Management Utility (PMU)

5

The Platform Management Utility is the operator interface to the Server Management system. The PMU is most often an application running on a PC 930 that is connected to the server management private LAN. It can also be an application running on the Multiprocessor computer system EV7 and accessing through the private LAN.

10

The PMU issues commands and requests as:

Server Management Protocol Packets to the PMU Server (see section 0 “PMU Server”),  
TELNET protocol to TELNET servers (see section 0 “Virtual Console Terminals ”)  
TFTP client protocol to TFTP servers (see section 13 “Firmware Load and Upgrade”)

15

Some characteristics of a PMU are:

A PMU runs on a trusted host on the private LAN. The PMU runs on a host that is a member of the SM LAN, and importantly, does not allow access from the customer LAN.  
This host connects to the private LAN directly or through a gateway.

20

Multiple PMUs can be operating simultaneously.

The IP address is obtained on the private LAN interface via DHCP. Two exceptions to this are a PMU running on an EV7, which does not use DHCP but uses the IP address assigned to the EV7 primary, and one reserved address for VMS.

25

A PMU interactively allows the operator to power segments of the machine, create and destroy partitions, query components and environmental.

A PMU enforces rules and policies about configurations.

A PMU implements policies reacting to error notifications.

Etc.

The PMU can also be an application running on EV7, and connecting through the CMM ram mechanism. Through this access mechanism, in an exemplary embodiment of the invention, 5 not all PMU functionality is supported. The CMM ram access mechanism does not support higher level networking protocols like TFTP and Telnet. Conversely, some capabilities of the SM protocol are provided only for the operating system, and not the PMU.

10 The summary table shown in Fig. 22 of SM Protocol packets issued by the PMU is presented here as the requests seen at the PMU Server.

Live configuration change means adding or deleting a cpu/io from a running partition  
Writing SRM environmental VARS and storing PCI slot information is only an OS callback action.  
This action is never initiated by the PMU.

15 A block diagram of the PMU Server is given in Fig. 23.

### PMU Server

20 The PMU Server is the focal point for PMU activity. It is available to the PMU at a fixed IP address, for example 10.253.0.0 which allows it to run on any MBM, and transparently fail over. The PMU Server has a dual functionality of (1) processing requests from PMU Applications and (2) distributing alerts to all PMUs. The PMU server simplifies the PMU interface by allowing the PMU to address any node on the SM LAN using the SM packet's common message header destination but still placing the PMU Server as the UDP datagram destination.

25

### PMU Server Handling of SM Requests

All requests from the PMU must go to the PMU Server. This allows the PMU Server to order commands if necessary, preempt operations when required, and synchronize operations.

Some PMU requests are handled right at the Server. The PMU Server caches some information about the system which it serves up directly to the PMU. These are “direct” commands.

Other times the Server acts as a command router just passing commands onto to the command target. These are “forwarded” commands. Some commands are forwarded to a Partition Coordinator (described in the next section) where a series of transactions can result. When forwarding commands to the partition coordinator, the PMU Server modifies the destination field to the IP address of the Partition Coordinator.

Sometimes, the PMU Server turns a single PMU command into a series of smaller tasks, masking the intricacies of the Multiprocessor computer system operations. These are “complex” commands.

The table of Fig. 24 lists the commands that the PMU Server services by group, class (forward, direct or complex) and a description of how the handling is done.

#### SM Commands Originating from the PMU Server

Only a few of the SM protocol packets actually originate from the PMU Server. They are generated as either part of power up system discovery, to satisfy cable testing requests, or to send alerts to the PMU. The table of Fig. 25 gives a list of SM commands originating from the PMU server.

#### PMU Server Handling of Alerts

When a micro needs to alert the operator of an error condition or noteworthy event, the micro sends the alert to the PMU Server. The PMU Server knows of all the PMUs that are connected to the system, and sends an unsolicited error message to all PMUs. The PMU Server distributes the alert to the LAN PMUs. Since the PMUs that are running on the EV7s through the shared ram interface are not known to DHCP, those PMUs are not sent alerts. The fixed PMU address range, used for operating systems that do not support DHCP, receive the alerts even if they do not exist.

The PMU Server distributes system events (sysevents) that span partition boundaries. Sysevents are notifications sent to the software running on the EV7. Example sysevents include (i) an environmental red zone, or (ii) a configuration change. More detail on the PMU server handling alerts is given herein below.

5

### Show Configuration with FRU Data Example

For an application program to display the Multiprocessor computer system Configuration, the following commands are used:

10 A series of REQUEST SYSTEM TOPOLOGY commands incrementing the entity number and saving the response data in an application structure hierarchy. When that is complete, MBMs, PBMs, CMMs, EV7s and power trays are known along with their IP addresses or rack number in the case of power trays.

15 All command requests from the application are sent to the PMU Server but using the IP address associated with the device being addressed (i.e. EV7, CMM, MBM, PBM). The application can make use of the commands:

20 GET MBM CONFIGURATION gets information about CMMs and RIMMs  
GET PBM CONFIGURATION gets information about IO7s  
GET MBM IP CABLING and GET PBM IO CABLING to display the Cabling layout of the Server.

25 GET PARTITION DATABASE to determine the current partitioning configurations.  
GET VOLTAGE READING on a component for the nominal and current values of all voltages.  
GET TEMPERATURE READING on a component for the limits and current reading of all thermal sensors.  
GET FAN RPM SPEED for all fans on component.  
GET POWER STATE to determine the state of all power supplies attached to an MBM or PBM.

5           GET EEROM DATA to identify the serialization information on sub-components that contain EEROMs. The EEROM Locator number used in this request is fixed to a given device type (e.g. for a CMM, 0=EV7 EEROM; 1-20=RIMM EEROM). In the case of an IO7 Drawer with a partition assignment, a GET PCI SLOT INFO request returns PCI Configuration Data for each slot in the drawer.

In Fig. 26 there is shown an exemplary simple configuration of 1 MBM, 1 CPU, 1 PBM, and 1 PMU is shown. The flow tables of Fig. 27, Fig. 28, and Fig. 29 reflect these hardware components. The MBM has the DHCP task, the PMU Server task, and the generic MBM protocol server task. These tasks are emphasized in capitals when they are involved in the flow.

10          Note that, in this flow, the EV7s are not affected. Information about the EV7s and RIMMs was gathered at "power on" and saved.

The MBM's CLI on the debug console replaces the PMU for the 2P Multiprocessor computer system Systems and is available on other configurations as well. It makes use of a set of commands to display the System configuration. A complete example for a simple Multiprocessor computer system configuration is shown in Fig. 30.

#### Partition Coordinator

The Group Leader starts a task, the Partition Coordinator 31,002 of Fig. 31, for each hard partition that is defined in the partition database. The Partition Coordinator has the overview of the state of all elements in its partition. It oversees and coordinates the operations between the elements. Routing and coordination of diagnostics must be done at a hard partition level, and is not done at the subpartition level. Therefore, there is only one Partition Coordinator per hard partition. The subpartitions in the hard partition are all managed by the same hard partition coordinator.

25

Since an MBM could be associated with as many as 8 partitions, an MBM could have as many as 8 Partition Coordinators running on it. The Partition Coordinator carries out the steps to start a partition, add and delete elements of a partition, and other operations that require syn-

chronization among the partition components. For example, during diagnostics, all EV7s of a partition are asked to send on their west port and receive on their east port. The Partition Coordinator directs all the EV7s to start this test at the same time. For another example, when adding an EV7 to a running hard partition, all EV7s in the partition must be quiesced, all subpartitions must be quiesced, and their router tables reconfigured. The Partition Coordinator directs all EV7s to quiesce at the same time, and it passes out the router tables to each EV7, and then continues all the subpartitions.

The Partition Coordinator is also the task that computes the routing tables for each EV7 in its partition. This compute intensive operation is broken up naturally on partition boundaries and can use the distributed MBMs as compute resources.

All requests and changes of status that take place on a partition are passed to the Partition Coordinator. The Partition Coordinator is instantiated by an MBM when a DISTRIBUTE PARTITION DATABASE message is received. Any changes in the state or attributes of the partition (i.e. all completed requests indicated below) cause a volatile database distribution message to be sent on a Train Broadcast.

#### Partition Coordinator Handling of SM Protocol Requests

The Partition Coordinator receives most of its requests from the PMU Server as a result of a user action. Most requests can be categorized into either of two types: (i) partition configuration and state directives, and (ii) partition data assignments. The configuration directives are usually “complex” requests that the Partition Coordinator turns into a series of simpler “direct” requests to each member of the partition. The partition data assignments are usually just redistributed by the Partition Coordinator and do not require acting on CMMs or EV7s.

25

The table of Fig. 32 reflects the actions taken by the Partition Coordinator on each partition request depending upon the current state of the partition.

## Protocol operations that Originate from the Partition Coordinator

The requests to assign data and parameters into the partition database are stored on every MBM/PBM. The Partition Coordinator embeds requests into a Train Full message for guaranteed delivery to each MBM/PBM. The configuration directives received at the Partition Coordinator are fanned out to all members of the partition. Sometimes these requests are turned into a series of actions, as can be seen in the flows referenced in the table "Partition Coordinator Commands Issued" of Fig. 33.

### Partition States

When a partition is started, testing is done on the IP and IO cable connectivity and the EV7s perform tests. With the results of these tests, the Partition Coordinator can calculate the EV7 routing tables. Routing a set of EV7s may fail. In that case the Partition Coordinator selects a subset of the EV7s and tries to route that set. When the routing is successful, the tables are distributed to the EV7s and the partition is activated.

For a configuration change to a running partition (add or delete), the partition must be quiesced, and goes back to the CONFIG ROUTING phase where new EV7 router tables are calculated. A PARTITION STATE DIAGRAM is shown in Fig. 34.

### Determining the routing table

Contents of an EV7's RBOX and CBOX IPRs determine its routing behavior within the mesh. The partition coordinator generates this information from the user description of the partition and routing options as well as from the actual topology of the partition as discovered by EV7 diagnostics, which are orchestrated by the partition coordinator. Inputs and outputs of the router algorithm are shown in Fig. 35.

### Connectivity test

Initially, the partition coordinator sets EV7 PID values and routing table entries sufficient for each processor to test its connections to adjacent processors as determined by the partition database and MBM cabling. Failing processors and connections are removed from the partition database and reported to the PMU.

## EV7 Routing

Within a partition each pair of processors must typically be able to communicate along a “dimension-ordered” path. It may also be possible to establish an “adaptive” connection. Depending on user choices and system defaults communicated from the PMU, “initial hop” and “SRC” routing may be permitted as well. Within the given constraints on the various forms of routing, the CTABLE module of the partition coordinator determines whether the partition can be routed. If it can be routed, CTABLE supplies routing table entries for each processor. Otherwise, failure to route is communicated to the PMU.

The routing process should always result in either a (1) success – the requested configuration is routed, (2) partial – only some of the EV7s in the partition can be routed, (3) fail – there are no good EV7s. A Routing glossary is shown in Fig. 36.

## Striping

Four processors can be grouped for balanced memory access (striping). When the partition database designates stripe sets, the partition coordinator creates them by assigning to the grouped processors PIDs which differ only in their two least significant bits.

## Clumping (EV7 SharedInvalBroadcast message)

Processors can be grouped to receive and forward invalidate messages (clumping). When the partition database indicates a maximum partition size greater than 20, the partition coordinator creates clumps by filling in routing table entries reserved for this purpose on each EV7.

## PID determination

The partition coordinator assigns a unique PID to each EV7 in the partition. PIDs are used to group processors for purposes of receiving invalidate messages (clumping) and of managing memory latency (striping).

25

The description of a partition in the database must include the maximum number of EV7s that will ever be members of the partition. Using this datum, the partition coordinator will determine if clumping is needed and the size of the clumps.

In response to the requests "Add EV7s to Running Partition" or "Delete EV7s from Running Partition", the partition coordinator analyzes the routability of the proposed new partition and proceeds if routing is possible. Otherwise, it reports an error. The partition coordinator will 5 also report an error if the Delete would remove from the partition some but not all of an active stripe set.

Active means that the members of the set are configured to access a non-empty set of memory by the striping mechanism.

### The Partition Coordinator Flow

10 The flow described in Fig. 37 is a high level flow of the important duties of the Partition Coordinator task in bringing up a partition.

At block 37,001, a task, the Partition Coordinator, gets created for each hard partition defined by the partition database.

15 At block 37,002, prior to this, all MBMs/PBMs have synchronized their partition data, and information about the current state of the system. With this information, the Partition Coordinator knows if the OS is running on the partition it commands. If the MBM was hot swapped or reset, it could be reinitializing under an already running operating system, which it must not disturb.

20 At block 37,003, if no operating system or SRM console is running, the new Partition Coordinator brings the partition into a known reset state.

At block 37,004, PIDs are assigned to each EV7.

At block 37,005, diagnostics are now run on each CPU in the partition. Some diagnostics require synchronized steps with other CPUs, and the Partition Coordinator coordinates that synchronization.

25 At block 37,006, the Partition Coordinator gathers all the results from the diagnostics run on each CPU.

At block 37,007 the Partition Coordinator uses the collected results of tests to configure the routing parameters.

At block 37,008 this process is repeated for several phases of diagnostics.

At block 37,009, determine which EV7 will be the primary cpu for the each sub-partition

At block 37,010, start the SRM on the primary, and start the secondaries. This is done for each subpartition.

At block 37,011, the Partition Coordinator services commands from the PMU which stop, halt, or in any way change the partition.

10 Configuration Events

This section discusses some of the configuration changes that occur on a Multiprocessor computer system. There is emphasis on the flow diagrams to show the interaction of the micros and the application of the system management (SM) protocol. Configuration events can be (i) soft configuration events (i.e. a reallocation of known resources), or (ii) hard configuration events (i.e. physical changes such as adding new hardware to the system.)

Some soft configuration events:

Creating a new partition

Starting a partition

20 Adding an EV7 to an existing and running partition

Adding an IO drawer to a partition

Destroying a partition

Deleting an EV7 from a partition

Deleting an IO drawer from a partition

Some hard configuration events:

Powering ON/OFF a partition

5 Hot Swap of a CPU module

Adding on new hardware, such as the addition of an 8P box

Creating a new partition

By default, the entire system is a single hard partition. The user can subdivide this system by partitioning the system. Partitioning the system requires creating an identifier for the partition and associating EV7s, IO busses, and memory with that partition.

10 Partitioning configuration guidance is offered by the PMU.

15 A hard partition always has at least one sub-partition, the free pool (perhaps called the “idle asset” pool) that is created.

In an alternative embodiment of the invention, memory can be shared by setting up a community.

20 As shown in Fig. 38, flows for creating hard partitions in response to operator action are shown.

Turning now to Fig. 39, Flows for creating a sub partition in response to operator action are shown.

25 Turning now to Fig. 40, A first part of a flow diagram for creating a new partition is shown.

Turning now to Fig. 41, the second part of the flow diagram of Fig. 40 is shown.

## Starting a Partition

This is an example showing the operation of starting a partition from an unknown state. The partition is already powered on. This flow is initiated by a user directive to reset the partition.

5

The "Reset State" is shown in Fig. 42.

The Diagnostic State is shown in Fig. 43.

A flow diagram showing the Configure Router State is shown in Fig. 44.

A flow diagram showing the Continue Partition State and Partition Running State is shown in Fig. 45.

## ADD vs. Move

Note that the MOVE EV7 TO PARTITION packet is used in the flow "Creating a New Partition Flow Diagram" of Fig. 40 because the partition is not actively running an operating system. In the flow shown in Fig. 46A "Add EV7 (Part 1)", an operating system is running in the partition and is the initiator of the configuration change.

In the Fig. 47, the use of "Add" is contrasted with the use of "Move". Hard Partition 0 / Sub Partition 0 is running an operating system. Therefore, any additions from the global free pool require the use of an Add. Within the partition, an Add is required to move an EV7 from the partition free pool into the running subpartition. However, a Move is used with the subpartition that is "Not running".

In contrast, a Move is used in all cases where the subpartitions are idle, and no operating system is disrupted by a configuration change, as shown in hard partition 1 (HPI) block of Fig. 47.

## Adding a CPU to a partition

Turning to Fig. 46A, there is shown an example of moving a CPU module (EV7#4 and EV7#5) from the free pool (hard partition #255 subpartition #255 ) to hard partition #2, subpartition #1. This flow starts after the operating system has been directed to perform the partitioning change. The flow continues in Fig. 46B.

#### Destroying a partition

A partition must be in Halt or Power Off state. Fig. 48 shows destroying a hard partition.

If no more partitions in Hard Partition, the whole partition is destroyed, including the SRM environment variables and any other partition specific non-volatile information.

Fig. 49 shows destroying a soft partition.

#### Deleting a CPU from a partition

The process of deleting a CPU from a partition has many similarities to adding a CPU. The partition must be quiesced and routability must be maintained.

#### Adding/Deleting an I/O drawer to a partition

IO7s migrate across hard partitions with the EV7s on their North port. Within a hard partition, IO7 can migrate freely without the need to quiesce the operating system. The role of the server manager in IO7 migration within a partition is minimal.

#### Power Management of a Partition

When a request is made to turn a partition on or off, the Partition Coordinator determines the requirements of each component as needed by other partitions. Power controls, powering components on or off, are performed according to component dependencies. EV7s depend on the VRM of the EV7 pair found on the CPU module as well as the 48V n+1 power supplies, and IO7s depend on IO drawer power supplies.

#### Power Off Partition

When receiving a POWER OFF PARTITION request, the Partition Coordinator should:

1. Send the POWER OFF command to CMMs where either both EV7s are assigned to the partition or the attribute state of the other partition is powered off or in the free pool.

2. Determine from the partition database whether any given PBM has IO7s completely under this partition's control or the state of the other partition is powered off or the other IO7s are in the free pool. When these conditions are met send the POWER OFF command to the PBM.
3. For all MBMs having at least 1 EV7 assigned to the partition, make a GET MBM CONFIGURATION request to determine the state of the 8 EV7s controlled by the MBM. If the state indicates it is not in use (e.g. being tested while in the free pool) and any other partition owners of EV7s controlled by this MBM are powered off, the Partition Coordinator may send the POWER OFF command to the MBM who turns off the power supply for the MBM.

#### Power On Partition

When receiving a POWER ON PARTITION request, the Partition Coordinator should:

1. Check, via GET POWER STATE, the power supplies on each MBM having an EV7 in the partition and send POWER ON to the MBM when the power supply is off.
2. Check, via GET POWER STATE, the power supplies for each PBM having an IO7 that is assigned to the partition and send POWER ON to the PBM when the power supply is off.
3. Send POWER ON to each EV7 in the partition
4. Continue bringing up the partition in the same form as reset.

#### Hot Swap of a CPU module (EV7 failure and replacement)

The flow shown in Fig. 50 and 51 is an example of SM protocol activity when a CPU module fails. The entire configuration is one partition, and the failing CPU module is not the primary.

The operating system does crash, and is restarted. The failed CPU module is replaced and the original complete configuration is restored.

Block 50,001 shows: CPU module is removed. EV7 and CMM are unavailable.

Block 50,002 shows: MBM polls CMM with GET CMM STATE. Timeout indicates the problem to the MBM.

Block 50,003 shows: MBM sends a SYSEVENT to primary EV7. Information is included to indicate a configuration change.

Block 50,004 shows: MBM sends an ERROR REPORT alert to the PMU server.

Block 50,005 shows: PMU Server redistributes the Alert to all PMUs.

Block 50,006 shows: The OS is likely to crash or hang in this situation. Consider the hang case. The CMM OS Watchdog fails.

Block 50,007 shows: CMM send an ERROR REPORT to parent MBM.

Block 50,008 shows: The parent MBM sends an ERROR REPORT alert to the PMU.

Block 50,009 shows: The parent MBM sends a SYSEVENT to the primary EV7. It has timed out so this SYSEVENT gets dropped.

Block 50,010 shows: The Partition Coordinator re-orchestrates bringing up the partition without the failed EV. This multistep operation is not shown.

Block 50,011 shows: The OS is running again on the degraded system. The CPU module is replaced. The MBM is polling with a REQUEST CMM STATE and discovers the new CMM.

Block 50,012 shows: MBM sends an ALERT to the PMU server.

Block 50,013 shows: PMU Server redistributes the Alert to all PMUs.

Block 50,014 shows: MBM sends a SYSEVENT to primary EV7. Information is included to indicate a configuration change.

Block 50,015 shows: The OS does REQUEST COMPLETE LAN TOPOLOGY requests, and discovers that there are new EV7s available.

Block 50,016 shows: The OS issues an ADD EV7 to RUNNING PARTITION to the PMU SERVER.

Block 50,017 shows: The Partition Coordinator re-orchestrates bringing up the partition with the new EVs. This multistep operation is not shown.

Adding a new 8P box to a system – “The Rogue”

When a new 8P box is connected to an existing Multiprocessor computer system, the operator must declare the new MBM as a valid member of the network. From the PMU, the op-

erator issues a command that includes the new 8P. All resources on the new 8P are entered into the free pool partition. Turning now to Fig. 52, a set Membership Flow Diagram is shown.

### Cable Configuration and Testing

A volatile copy of the cable connections for both IP and IO cabling is prepared by the PMU Server after a group is formed. PMUs and Partition Coordinators can request the RECONFIGURE CABLING to have the PMU Server re-check cable connections. 8 IO ports and 6 IP ports are located on each MBM and 4 IO ports on each PBM. Each cable's connector contains: 1) an input signal to determine if the cable is present, 2) two sets of LEDs marked A & B at each end of the connection, 3) the ability to turn any set of LEDs on or off and read the result. A LED lights on both ends of the cable when zero is written to the cable test signal. To conserve power, LEDs should normally be extinguished.

### IP Cable Discovery

As part of the determination of the proper inter-processor physical cabling layout, each MBM cooperates in the PMU Server's cable testing by responding to the requests for GET MBM IP CABLING. IP cables connect pairs of EV7 external 8P routing ports. The 6 IP ports are 2 North, 2 South, 1 East and 1 West.

An IP cable configuration diagram is shown is Fig. 53.

The MBM requests a RECEIVE CABLE ID from each MBM expected to be on the other end of each of his external ports. The MBM sends the ID down the cable by performing the equivalent to SEND CABLE ID request on his own ports. Whoever responds positively to the request returns the receiving ends port direction and EV7 port number and the EV7 ID received. This is then repeated for all external port directions on the MBM. Each MBM known to be present in the Multiprocessor computer system reports his cabling connections to the PMU Server who maintains the cabling database.

There are certain illegal thumb-wheel number settings for rack thumb-wheel and MBM thumb-wheels as well as invalid cabling configurations that are checked at this time and reported as an error in the error log. The MBM thumb-wheel number consists of a rack number (0-7) and an MBM number (0-3). Turning now to Fig. 54 coordinate addressing relationship to thumb-wheel addressing is shown.

## IO Cable Discovery

The PMU Server makes a similar determination of all IO7 to EV7 cable connections by making a GET PBM IO CABLING request to each PBM. The PBM in turn sends a RECEIVE CABLE ID to all MBMs; then sends the IO7 riser ID down the I/O cable via a SEND CABLE ID request. The PBM responds with the EV7 IDs connected to the I/O risers. The PMU Server repeats this for all PBMs while maintaining the cabling database, as shown in Fig. 55.

## Cable Database Retrieval

Retrieval of the cable connection database is done through requests to the PMU Server. When the PMU or Partition Coordinator requests the cabling via GET CABLING CONFIGURATION, the results from the last RECONFIGURE CABLING are returned.

A GET CABLE CONFIGURATION Block Diagram is shown in Fig. 56.

## Cabling Check on Additions

The partition coordinator uses the cabling database to determine the connections that can be made among the processors assigned to the partition under its control. When adding or removing an EV7 or IO7 from/to the partition, the partition coordinator must determine the effect this change has on its own grid of processors. The result may leave other processors isolated when attempting to route. The PMU is warned of the resulting effect. Fig. 57 shows a cable Addition/Deletion Block Diagram.

## PMU Cable Locator Assistant

To assist the cabling operation itself, the PMU makes use of SET CABLE TEST SIGNAL STATE, GET CABLE TEST SIGNAL STATE and SET ATTENTION LIGHT INDICATOR. A block diagram of a Cabling Assistant is shown in Fig. 58.

This could be used to identify the proper connection by writing 0 to the cable test signal on the port of a cable that is being tested. Then read the same port on the expected other end of the cable to determine if the read of the signal is present. The MBM rack, PBM Rack or CMM attention light indicator can be lit or flash to identify the location until the cable is properly con-

nected and thereby assist in locating the other end of the cable. When the cable is properly connected, all LEDs should be extinguished to conserve power.

### The MBM / CMM / EV Hierarchical Relationship

An MBM is parent to as many as 4 CMMs, and grand parent to as many as 8 EV7s. Communication to the CMMs is over a serial connection running PPP. The MBM provides some important services to both CMM and EV7. Services provided by the MBM comprise:

- Proxy Forwarding
- CMM presence watchdog
- TFTP Server
- Error repository & distribution
- Power hierarchy
- Time Services
- Virtual Console Terminal
- SRM Env. Var. Repository

### Proxy Forwarding

Turning now to Fig. 59, a block diagram of Proxy Forwarding is shown.

The MBM connects to the LAN, but the CMM does not. The MBM enables proxy arp for each of the 4 PPP links. It configures host routes for each of the links. Packons on the LAN destined for the CMM are received by the MBM and forwarded within the IP stack to the appropriate CMM.

The EV7s are never UDP sources or destinations, so forwarding is only provided for the CMMs.

The MBM also provides the forwarding of CMM request to the LAN.

The PPP interface supports UDP, TFTP, PING, and the SM protocol, and does not support TCP or Telnet. SM private LAN broadcasts are not sent to the CMM.

### CMM Status Check

The initial PPP connection to each of the 4 CMM to MBM serial lines brings up the connection and sets up the proxy arp entry so that external traffic gets forwarded to CMMs. The

MBM continues to check the state of the CMM connection and attempt to re-establish with any that are not functioning. This service also requests the state of the CMM and its EV7 on a periodic basis of about 2 seconds. Timeouts imply the connection is lost, in which case, an error entry is made and attempts are made to re-establish the connection at periodic intervals.

There are several ways for an MBM to detect if a CMM has stopped functioning:

1. The PPP layer can detect if the PPP connection is lost.
2. A hardware register indicates the presence of a CPU module.
3. An SM protocol packet could fail.

### Virtual Console Terminals

Turning now to Fig. 60, a block diagram of the virtual terminal overview is shown.

Each MBM has one Telnet server for each EV7 virtual COM port (8 total for COM1 ports, 8 more for future use of COM2).

#### Addressing the virtual terminal

Each Telnet server has the MBM's IP address, and a unique port number, starting with port 323. RFC 1161 shows 256-343 as unused in the Well Known Port list.

The Port numbering scheme is shown in Fig. 61.

#### Establishing a Telnet session to a partition

Every EV7 has a virtual COM1 and COM2 port. However, the server management firmware only provides console terminal support for the primary EV7 in the partition.

#### Steps to establishing a connection

1. Connect PMU platform (e.g. laptop) to an available hub port.
2. Get the IP address/port for the Telnet server for a Partition's Primary EV7 (PMU issues GET TELNET IP ADDRESS/PORT to PMU server. PMU server responds with Telnet server's IP address and port).

3. Telnet directly from PMU to Primary EV7's telnet server on Primary EV7's MBM.

Note: PMU connects directly to Primary EV7 Telnet server (not through the PMU server).

Flow diagram of a virtual console session

A Virtual Terminal Flow Diagram is shown in Fig. 62.

Telnet server technical details

The telnet server provided by vxWorks will not do exactly what is needed for this design. The vxWorks telnet source code must be modified to allow more than one server, and to telnet to a task other than the target shell. The telnet server will take the keystrokes, and issue the `PUT_CHAR` commands to the virtual console uart on the CMM. Each telnet server running on the MBM sets up for:

- WILL ECHO
- DO SUPPRESS GO AHEAD
- WILL SUPPRESS GO AHEAD
- DO TERMINAL TYPE
- DO TRANSMIT BINARY
- WILL TRANSMIT BINARY

Noteworthy virtual console behavior

The OS and SRM console always communicate through the primary EV7. The primary is capable of changing (e.g. A `SWITCH PRIMARY` EV7 command occurs) while a telnet session is running on the PMU. The MBM breaks the Telnet connection when the switch primary occurs. Upon recognizing the broken connection, the PMU reissues the `GET TELNET IP ADDRESS/PORT` and reestablishes the connection to the new primary.

## VIRTUAL TOY CLOCK

The virtual time of year clock (TOY) clock is next described.

Turning now to Fig. 63, a flow diagram for SET BASE TIME is shown. In summary, all of the MBM units in the multiprocessor computer system maintain a "base time". The time as

seen by each partition, the partition "current time", is computed from the base time by adding a "delta time" stored in the replicated database on all MBMs and PBMs.

The delta time for each partition is initialized to zero upon boot up of a partition. When a partition is initialized and booted for the first time the EV7 processors accesses the time of day data via the CMM. The CMM computes the time for the partition by adding the partition specific delta time to the base time. If the operating system running on that partition changes the time of day, the CMM and the MBMs store the difference between the time value being set and the base time as the partition's "delta time".

The TOY (time of year) chip, also known as the "watch chip", on each MBM supports TOY functionality for the client CMMs and CPUs. Neither the CMM nor the CPUs have TOY hardware, so the MBM must supply the TOY information. The MBM periodically sends each client CMM a DISTRIBUTE BASE TIME CHANGE SM packet. The CMM applies the delta time, then stores the TOY data from this packet in address space that is accessible through the FPGA as BB\_WATCH. An EV7 reads BB\_WATCH through the Gport interface.

The Server Management strategy regarding TOY updates assumes that the operating systems access the TOY infrequently. Also the strategy assumes that  $\frac{1}{2}$  the update interval is within tolerance for the operating system.

The watch chip is initialized by the application program with the SET BASE TIME SM message.

### Synchronization of the Watch Chip

Each MBM contains a physical time of year device, a TOY clock 16,016 or a 19,024 TOY device, also referred to as the watch chip mentioned above. These time pieces are all synchronized to the same base time, which is an arbitrary time set by the user. The base time could be Universal Coordinated Time (UTC), Eastern Standard Time (EST), etc. Synchronization is maintained within a predetermined time envelope by use of a "synch wave" technique. The predetermined time envelope is determined by the known delays in messaging between the microprocessors of the MBMs and PBMs. The "Synch Wave" technique is described in the following paragraphs.

A hardware clock consists of an oscillator which generates a cyclic waveform at a uniform rate. A counting register records the number of cycles elapsed since the start of the clock. For example, a quartz crystal controlled clock may be used for the oscillator. Quartz clocks of standard engineering design may have a maximum drift from real time, of a drift rate within a microsecond per second. When a number of processors exchange messages to coordinate their clocks, the delay in message transmission gives an upper bound to the error between readings of the various clocks of the processors.

The clocks of various MBM microprocessors diffuse messages carrying their time determination by a microprocessor sending a new synch message to all other processors. All microprocessors receiving a new synch message transmit their own synch message to all other microprocessors. This process generates a "synch wave" as the message diffuses to all microprocessors. In a preferred embodiment of the invention, the fastest microprocessor clock may be used to define the base time. One clock will be a little faster than any other clock. The microprocessors all set their clocks to agree with the fastest clock setting in the synch wave. All clocks are then in synchronization within a time envelope determined by delays in message propagation between processors, and all are set to the time of the fastest MBM microprocessor clock. The synchronization of all MBM and PBM microprocessors by the synch wave method produces the base time for the partition of the MBM and PBM microprocessors.

In an alternative embodiment of the invention, the base time is set as follows. The MBMs (multiprocessor system wide) elect a principal MBM unit. The MBM units generate synch waves through which they determine the fastest MBM watch chip clock. The principal MBM unit adopts the fastest time as base time, and sends messages to all other MBMs telling them to write the adopted time as the base time into their watch chips. However, this method is not preferred because some clocks in a partition with a fast clock may be reset backwards, and such an event could have collateral consequences.

An alternative method for setting the base time is for the principal MBM, after its election, to simply use its watch time as base time, and eliminate synchronizing to the fastest clock in the MBMs of the multiprocessor computer system.

When a new MBM enters a system, it must recognize that its base time is invalid. The new MBM must then get the correct time from the next “sync wave”. There is a base-time synchronizer on each MBM.

The SET BASE TIME command is executed by an outside agent, usually a human. The SET BASE TIME command provides the MBM group leader with a time to use as the current base time. The MBM group leader time then becomes the new real current time throughout the entire multiprocessor computer system through the following process. The base-time synchronizer routine sends out a “synch message” with the new base time to all MBM's and PBM's. This base time is then written into all watch chips in the MBM and PBM units of the multiprocessor system.

The base-time synchronizer is awakened by one of the following events: a new base time is received (i.e. A user manually sets the time), a new processor joins, a new time message is sent by another MBM, or a timeout occurs (i.e. The base-time synchronizer may do periodic broadcasts of time).

Delta time is a partition specific value that is determined on each partition. The delta time is maintained on each MBM in a partition. The delta time plus base time equals BB\_WATCH time, which is needed by the EV7s. The CMM uses GET DELTA TIME to get the delta time, then combines this delta with the base time to provide BB\_WATCH time.

All EV7's must have direct access to a BB\_WATCH server. Therefore, there will be one BB\_WATCH server per MBM, supplying base and delta times to each CMM for the EV7s. Each BB\_WATCH server is kept in synchronization with the other BB\_WATCH servers via the base-time synchronizer function.

See the following paper for more details of sync wave synchronization of a group of processors such as the microprocessors of the MBMs and PBMs: “Clock Synchronization in the Presence of Omission and Performance Failures, and Processor Joins”, Flaviu Cristian, Houtan Aghili, and Ray Strong, IBM Research, Almaden Research Center, 1986, all disclosures of which are incorporated herein by reference.

The delta time for each partition is maintained in the replicated data base 84,000 shown in Fig. 84A and Fig. 84B, in block 84,008. Through the replicated database each MBM and PBM unit knows the delta time for each partition. Further, the base time is the same for all

MBM and PBM units. Accordingly, any partition may be broken down and reassembled using any of the physical units of the entire multiprocessor computer system, and have the correct current time available for the new partition's EV7 processors. The correct current time for the new partition is taken from the common base time and the delta time for that partition which is read from block 84,008 of the replicated database 84,000.

## DHCP Server

The DHCP Server supported by VxWorks is started by the Group Leader in order to give lease IP addresses to PMU DHCP clients that connect into the SM LAN. A DHCP block diagram is shown in Fig. 64.

Any DHCP leases made cause a DISTRIBUTE DHCP LEASE DATA request with the IP Addresses, MAC Addresses and Lease Duration information. This request is sent to all MBM/PBM peers, so each can maintain their copy of DHCP leases in NVRAM.

The DHCP lease data is maintained in order to restart the DHCP Server with the proper lease data whenever the group leader restarts it, and so the PMU Server can send alerts to all known PMU IP Addresses on the SM private LAN.

The DHCP Server is started by calling dhcpsInit() and in turn controlled through the following customized pieces:

- dhcpsLeaseTbl structure that contains the entries available to lease and their associated lease parameters. In our case we only need a single entry consisting of:

```
{“dflt”, “10.253.0.1”, “10.253.0.253”, “snmk=255.0.0.0”, “maxl=604800”};
```

This uses the default name entry with an IP address range of 10.253.0.1-253 using a subnet mask of 255.0.0.0 and a maximum lease time of 1 week in seconds.

- The hook routine `dhcpsLeaseStorageHook` that is invoked by the server with an action that the hook routine should take to store or retrieve the current DHCP hashed bindings database settings, which in our case will be copied to/from the MBM Nvram. To interpret the format of these binding database entries we can make use of the DHCP Library hash functions. In the case of a storage request, the copy is sent on the Reliable Database Train as a Distribute DHCP Lease Data request.
- Since the vxWorks DHCP Server does not inform us of database changes in a timely manner, the task “`own_dhcps`” is started to periodically obtain this information from the DHCP Server’s database.

### Non-Volatile Storage

The MBM flash contains the firmware images for all the micros, FSLs for all the micros, FPGA code, EV7 diagnostic firmware, and the SRM console. Additionally, the flash is used to store error logs, SRM environment variables for each partition, and all partition configuration information (referred to as the partition database).

### Flash Layout

The following table estimates the sizes of the data stored in the flash. Items with file names are accessible with TFTP for read and write. The MBM image of VxWorks kernel and libraries prior to adding application code is 0.7 MB uncompressed. The MBM Fail Safe Loader in compressed format is 0.35 MB. Since decompressing the MBM image would cause a slight delay in booting, it is not used for the MBM Image.

Fig. 65 is a table giving Flash Layout.

### Image Header Format

Turning now to Fig. 66, an Image Header is shown.

All images use the APU header format used on Alpha Systems. Some fields are set to fixed values for backward compatibility with existing tools. The header begins in the second 512

byte block of the image and the images make use of a 32 bit checksum in the last 4 bytes of the image.

A Flash Driver is required to perform the necessary reads and writes on the file descriptor argument in vxWorks “tftpGet”. The driver is opened using the TFTP filenames and an indicator of a test RAM version when requested. On write requests, the Driver is responsible for erasing flash sectors and backing up other files contained in the same sector. The driver read requests are simply memory copies from ROM.

### SRM environment variables

Turning now to Fig. 67, on SRM Environment VARS Flow Diagram is shown.

The SRM console requires non-volatile storage to maintain environment variables, such as the boot device. NVRAM is not available on the CPU module, so the SM subsystem makes available some of the NVRAM managed by the MBM. The SRM console, via the SM protocol, reads and writes its block of environment variables to the CMM. The MBM maintains a 2K block for SRM environment variables, backed by NVRAM, for each active partition. In order to maintain a partition’s environment variables even over partition component adds and deletes, the environment variable information is replicated on each MBM.

### Error cases

If all MBMs and PBMs are not accessible, and therefore replication cannot occur to all micros, then the environment variables are only stored in volatile memory. Only the partition primary CPU manipulates the environment variables, so stale information on inaccessible MBMs is not used.

### Flash Organization

The CDP Flash, which may be adopted in an illustrative embodiment of the invention, is organized in 32 bit wide accesses. Each word may be in a different device. Each device is 2MB with 64K sectors. Since it is 32 bit wide we can consider a logical sector as 256K and with 32 sectors available in 8MB.

## Firmware Upgrade and Test Load

Multiprocessor computer system firmware is upgraded from the PMU over the SM private LAN. A firmware load and upgrade block diagram is shown in Fig. 68.

### Upgrading Firmware

Turning now to Fig. 69, a flow diagram of upgrading CMM firmware is shown.

#### Upgrading MBM/PBM firmware

The PMU can upgrade individual MBM/PBM firmware images from files it supplies. The PMU Server issues an UPGRADE FIRMWARE command to the target MBM. The MBM opens up a TFTP session back to the PMU. The MBM pulls over the specified file from the PMU, writes its flash part, and notifies the PMU of success.

#### Upgrading CMM firmware

Upgrading CMM firmware is handled slightly differently than an MBM/PBM upgrade. The PMU issues the UPGRADE FIRMWARE command to the parent MBM specifying CMM firmware. The MBM updates its flash part. After its flash is updated, the MBM sends an UPGRADE FIRMWARE command to the first CMM. The CMM opens a TFTP session to the parent MBM and pulls over the specified file from the parent MBM's flash. The CMM writes the file to its flash part. The remaining CMMs are updated sequentially in the same way. The CMM only TFTPs to its parent MBM.

#### FSL recovery upgrade

##### FSL firmware upgrade of a CMM

On CMM startup, should the CMM FSL program detect a bad checksum on the CMM firmware, the FSL will start a TFTP session and pull over a new copy of the image from the parent MBM. This image automatically gets written into the flash.

##### FSL firmware upgrade of an MBM

On MBM startup, should the MBM FSL program detect a bad checksum of any of the images in the MBM flash, the FSL will bootp request, from its LAN peers, a new copy of the image. The

MBM starts a TFTP session and pulls the new image over, writes it into its flash, and resets itself and reloads.

## MBM OS Components

### VxWorks

Wind River's VxWorks has been selected in an exemplary embodiment of the invention as the RTOS of choice to support the MBM application. The VxWorks kernel has priority driven pre-emptive scheduling in a real-time multitasking environment with several choices of inter-task communication methods, interrupt handling support and memory management. Libraries are readily available in x86 compiled format to support our networking and serial device needs. The Wind River Tornado2 development environment contains a set of tools to build VxWorks bootable and downloadable images and perform source level debugging.

Additional libraries beyond the default set used for MBM application (based on the ElanSC520 boot extension of the 486 BSP) build are:

- END Interface Support for the END Network Driver
- PppLib for point to point connections with the 4 CMMs
- DHCP Server for maintaining PMU IP Addresses and Lease Times
- Telnet Server for starting SRM Console CLI interfaces with the PMU
- TFTP Client and Server for transferring images between CMM, MBM and PMU

### Boot Program

The boot process performs x86 specific processing to place the processor into 32bit flat protected mode and sets up the ELANSC520 HW registers to allow the vxWorks kernel to start.

### MBM Initialization

Following the rules defined in the VxWorks Programmer's Guide in the section on Board Support Package, we have made changes to initialize the Hardware Registers in sysLib.c. The function initElanRegs was added and called from sysHwInit and sysHwInit2. In addition PciAutoConfigLib was added, and minor modifications were made to sysNetConfig.c for supporting PCI Initialization.

On completion of the kernel initialization the usrAppInit entry point is programmed to perform the POST checks and initialize a list of services, as follows:

On 2P versions, initialize the CLI Handler to control the Debug COM Port.

Run the Power On Self-test on UARTs, Flash devices, I2C Master and slave device accessibility and on the NIC. An errors found are reported on the OCP and CLI port.

Some errors may be fatal. Some errors may have been detected in the kernel boot phase and error indicators set on the LEDs and UART for the CLI port. Where possible errors will be logged to the MBM Error Log and displayed on the OCP error line.

Initialize the error log to check on log validity, set up pointers, log any POST errors.

Initialize the ECC Error Handler to keep track of single bit ecc errors within the MBM SDRAM.

Determine by checking the switch or jumper whether we are in MBM or PBM mode.

Read the MBM & Rack thumb-wheels and start-up the network interface using the derived IP address for the MBM. Start-up 4 ppp connections to the CMMs to allow for automatic IP message forwarding to each CMM. The initial connections are established through the use of the vxWorks function, usrPPPInit, with arguments of the derived IP addresses of each CMM, and PPP options: PPP\_OPT\_PROXYARP and PPP\_OPT\_PASSIVE\_MODE.

Initialize a service to periodically check on the availability and health of each CMM.

Initialize a service that initially reads and caches the I2C device data and then periodically updates the dynamically changing portion of this data.

The following list of services is determined to be started in every MBM and have a distinct function of processing externally received messages. The services are separated so as to improve expediency and avoid latencies in processing requests that require device accessing:

Start the Group Message socket reception processing of LAN Group Messages as defined in the network protocol specification

Start the Atomic Update Message socket reception process to perform the Data Base Update services as defined in the network protocol specification

Initialize a service to process Partition Control Messages

Initialize a service to process Firmware Load and Upgrade requests using TFTP and start the TFTP Server.

Initialize a service to process Environmental/FRU Messages

Initialize a service to process all other commands, which currently seem to not require any device latency and therefore get bundled in this service.

Now that the list of request processing services is started; start the service that sends application requests on the SM LAN Command/Response socket and the service that receives on that same port.

The remaining services on the MBM are started via a request from the group leader to delegate a system wide service to the MBM.

The final service, MBM Watchdog, is started in a non-debug environment to insure the MBM stays alive.

#### MBM ECC Error Handling

The ELANSC520 SDRAM controller has provisions for correcting single bit ECC errors and detecting multi-bit errors. Enabling ECC error checking causes some latency in accessing SDRAM.

#### I2C Bus Driver

The I2C bus driver accepts requests to I2C devices on the I2C buses. The hardware device chosen for the I2C masters are Philips 8584 (See PCF8584 Data Sheet for details all disclosures of which are incorporated herein by reference). Internally the driver interacts with the 8584's Control/Status Register and Data Register to send the request to I2C device being addressed. Note: There are no multi-master requirements on the I2C buses. I2C buses beyond zero require a selection on the part of the driver setting the multiplexer register.

The vxWorks application interface to this driver consists of the following commands: open, ioctl, read, write and close. The I2C Driver also makes use of an ISR routine to capture LM80 interrupts caused by thermal or voltages being out of range. The following implementation note suggests a means of keeping the current offset for a device atomic across groups of reads and writes:

## I2C Request Handling

There are several requests that require accessing I2C data that can dynamically change (GET VOLTAGE READINGS, GET TEMPERATURE READINGS, GET FAN RPM READINGS, GET POWER STATE, GET SWITCH STATE, GET OCP SWITCH STATE & any request involving the thumb-wheel settings or LED values). These requests as well as others are made for environmental sensor readings found on I2C slave devices containing thermal, voltage, fan speed, digital switches. Since many sources can be requesting this data, avoidance of I2C bus latency is enforced by taking data from a global structure that is cached in RAM. On a periodic basis of approximately 2 seconds a refresh is performed on this cached data. The EEPROM device values are also cached and only read in at initialization time.

Write requests to I2C devices (SET FAN RPM SPEED, SET OCP DISPLAY DATA, SET ATTENTION LIGHT INDICATOR & SET EEPROM DATA) get written first to cache and then to the device itself via the I2C Bus Driver. During any write phase a global flag is set to insure that the refresh process synchronization does not overwrite the new data.

A fixed table of actions to take when sensors go out of range is also performed by this service. At initialization the LM80s are setup to interrupt on the Warning condition, which when reached is changed to the Failure Condition.

## MBM Watchdog

The ELANSC520 CPU has an internal watchdog reset capability that is used to ensure the MBM functioning despite firmware failure.

The MBM watchdog task makes use of the VxWorks timeout delay to periodically update the SC520 CPUs watchdog timer so that a reset is not triggered. If the MBM system hangs the task fails to update the watchdog and the SC520 Hardware causes reset to occur.

This MBM watchdog task's usage and timeout parameters are controlled by compile time parameters. The time out duration is in seconds to insure re-triggering the time out counter is not influenced by other task priorities. As part of the MBM initialization phase, the reason for reset is detected and written to the NVRAM error log and the PMU alerted. As with any other MBM

or PBM reset, group formation takes place and CMMs are checked for availability and the current state of each EV7.

### Failsafe Loader

The MBM-FSL Image load is invoked by the boot program, when either the MBM image is corrupt or the FSL jumper has been set on the MBM. The MBM FSL requires a subset of the standard application.

The failsafe loader is limited in capability. It's duties are:

- Initialize the error log by checking the log's validity and setting up its pointers.

- Initialize the CLI handler to control the Debug COM port.

- Invoke an error entry that will write a message to the OCP indicating a firmware upgrade is required, an error log message and a PMU alert.

- Read the MBM & rack thumb-wheel and start the network interface using the derived IP address for the MBM.

- Start the TFTP/Bootp request for the MBM Image and write the image packets to flash.

- On completion of this process, reset the MBM.

### Error Logging and Reporting

The MBMs devote a section of their non-volatile storage to error logs. Error such as environmental errors and network errors are stored for later retrieval. In addition, MBMs notify the PMU and partition primary EV7s of errors. A mechanism to allow PMUs to retrieve error logs is described here.

### Error Log Format

Turning now to Fig. 70, an Error Log Entry Format is shown.

A circular error log is maintained in a group of flash sectors. At initialization, the global pointers to the first and last entry are set as well as the current entry number. A flash sector is erased each time the log approaches the next sector, and the first entry pointer is recalculated. Error log entries are variable in length and use the format shown in Fig.

The end of log entry is identified by an entry number, timestamp and entry size value of 0xffffffffffffffffffff.

The error entry has a header and a data portion. Only errors that require special data values besides those found in the header require a data portion. A format for ERROR ENTRY DATA FORMAT is shown in Fig. 71.

The text portion returned from an error log entry retrieval is not kept in the error log but rather as a part of the MBM/PBM program. The Device Error Codes 1 and 2 are reserved for free form text errors and binary form text errors respectively.

### Error Log Controls

The PMU application can gather error log data from all MBM/PBMs by polling those entities for the GET ERROR LOG COUNT and obtaining any new entries via GET ERROR LOG ENTRY. When adding new MBM or PBM components, or fixing errors that were reported, the PMU or operator should decide whether the ERROR LOG CLEAR is to be used.

The PMU error log database, as a minimum, should consist of the error log count for each MBM and PBM in the Multiprocessor computer system. The format and contents of the PMU database is not specified in this document.

### Error Reporting Message Flow

When an error is reported by a Server Management entity, the actions to be taken in reporting the error depend on the severity. The system-wide table of "Error Messages and Actions" is used by the MBM, PBM and CMM firmware to determine the actions to take on an error.

As an example, we can follow the actions in reporting an EV7 thermal overheat condition by a CMM through the flow diagram of Fig. 72.

### FRU analysis from the SM subsystem

### General Diagnostic Flow

The Partition coordinator is responsible for bringing a hard partition into an initialized state. This involves sequencing the hardware through power-on, diagnostics, and SRM firmware load.

Diagnostics are performed, (and FRUs potentially indicted), whenever a hard partition is completely re-initialized. This occurs initially at power-on, but can also occur upon the creation of a new hard partition or upon an operator reset request. Individual CPUs are also initialized and re-diagnosed, when they are hot added to a running partition. At several points in the initialization process, the possibility exists that a failing response (or timeout) will occur. When this does occur, the partition coordinator must make a decision on what hardware capabilities to disable (or simply not make use of).

In the case of diagnostics in which there is interaction between components, this may require collating the results from multiple CPUs, running through a decision tree, and forming a new result.

#### FRU Examples

A module, cable, or assembly that can be removed and replaced in the field is a FRU. In some cases, for the purpose of diagnosis and repair, additional detail is made available to help isolate the failing component.

#### FRU Indictment

When a failure occurs, a field replaceable unit (FRU) must be identified and replaced. A decision tree is used to collate errors from multiple sources concerning a FRU. Each test includes instructions on the FRU callout determination. This applies mainly to the tests which involve the interconnect between CPUs where a decision must be made whether to indict the EV7s or their interconnecting cables and etch.

Diagnostics include not only the SROM and XSROM EV7 tests, but also the power supply, fan, and cable test status determined by the MBMs. For instance, a PCI drawer that contains no working power supplies would be mapped out of operation. An 8P backplane in which the fan has failed would not be used as part of a partition.

FRUs that fail completely must be marked as failed in the list of available resources to configure.

For instance, router configuration must not try to configure the RBOX of an EV7 that has experienced a fatal error.

#### FRU Callout

Referring now to Fig. 73, a block diagram of FRU logic is shown. Test results 73,002 are combined with "FRU DETERMINATION RULES" 73,004 by logic 73,006 to generate FRU Reports 73,008.

Upon completion of indictment the "Reported FRU" is then reported by the partition coordinator and logged using the Error Reporting packet. It is up to the PMU to display these errors as they are received. They should also be available to the operator by having the PMU interrogate the MBM/PBM error logs after POST has completed.

#### Partial Failure

In some cases, the partition coordinator will have to decide whether to make use of a FRU even when a test (or tests) have failed. In the case of interconnect failure, that link must not be used by the routing computation or enabled from either side. In the case of failing memory, a CPU may still be used without configuring the memory.

#### Unroutable Nodes

Due to the architecture of the EV7 router, the configuration of a partition which results from mapping the user's desired configuration onto the population of good CPU nodes may result in some nodes being unreachable by the router. The partition coordinator applies the appropriate heuristics (possibly as simple as forming the largest possible rectangle) in order to create a routable subset. This results in some good CPUs remaining unused by the initialized partition.

#### OCP Display Format

Assuming that the OCP has 20 columns and 4 rows, the layout in an exemplary embodiment of the invention for the display area depends on the state of the MBM 8P. Fig. 74 shows the display in table form. During the MBM initialization the format can be depicted as follows:  
Where,

Overall Progress – can take on the values:

“MBM PowerOnSelfTest”,  
 “MBM Initialization”,  
 “MBM Group Formation”,  
 “Server Management Ready”

Current State – takes on values of steps within the Overall Progress.

Location within state – takes on detailed values of the address or identifier within the current state.

Error Message – on failure, the message depicts the type of failure at the state and location identified.

When the partition(s) that the 8P MBM has been assigned to are started, the OCP display takes on the format shown in Fig. 75.

Where the data on each line is as follows:

1. The 1<sup>st</sup> line, the CPU line, has a simple EV7 number of the location within the 8P when present, a blank when not present and an X when failed.
2. The 2<sup>nd</sup> line, the memory line, is used to denote whether the memory test passed P or F failed.
3. The 3<sup>rd</sup> line, the power line, is used to denote presence of power 1 or absence 0.
4. The 4<sup>th</sup> line, the error message line otherwise contains indicators of the OCP switches located in the display as shown in Fig. 76.

The MBM firmware clears the error line when rechecking the state of the failing part.

### OCP Switch Control

When the MBM receives an OCP SWITCH ASSIGNMENT, it stores the value in its private copy of NVRAM, which is later used to determine the action to take when the switches are depressed. Fig. 77 shows a block diagram of an OCP switch process.

The MBM polls the switch status of its OCP buttons to determine if an operator has depressed the reset, power on/off or halt button. The current state of the OCP switches is maintained by the MBM for recall when a GET SWITCH STATE is requested. It is also used by the MBM to determine a change in the state of the switches. In this case, for EV7s physically located

on this MBM, the MBM takes action himself; otherwise, the MBM sends a reset, power on, power off, halt enable or halt disable command to the PMU Server for either the entire system, or for the partition assigned to the OCP. A determination of which action to take depends on the current OCP assignment for the MBM.

### Miscellaneous Command Handling

There is a group of commands that are bundled into “miscellaneous,” and also a set that have not been covered above. The table of Fig. 78 describes the processing methods on these commands:

### CLI Commands

A serial communications port is dedicated to the processing of a command line interface. Commands entered on this line are forwarded to the PMU Server and used to depict the configuration . This command line interface can only handle commands used to control the 2P Multiprocessor computer system Platform, which only has the ability to run a single hard partition. The commands available at this interface are shown in the table of Fig. 79.

### MODEM Operations

#### Modem Control on CLI Port

The Multiprocessor computer system 2P configuration makes use of the knob settings shown in the table of Fig. 80 to establish a modem to modem connection on the CLI serial port:

### Reliability, Robustness, and Resilience

Failures in the SM subsystem can affect the Multiprocessor computer system. Once a partition is up and running, the role of an MBM is limited and probably not catastrophic. A partition that is not up and running might not start, depending on the configuration, with a failed MBM member. The impact of the different failures that could happen are described in this section.

## MBM Failures

### MBM failures on an idle partition

The effects of an MBM failure in an idle partition vary depending on the configuration. If there is only one MBM in the partition's configuration, the partition is not capable of being started. If there is more than one MBM in the partition, that suggests that EV7s behind each of the MBMs are in the partition. Therefore, if one of the MBMs fails, the partition may be able to start in a degraded configuration without the resources behind the failed MBM.

### MBM failure effects on the CMM in a running partition

Although important, many of the services of the MBM to the CMMs are not critical to a stable, running partition. The effects of an MBM failure to its CMM children are described in the table of Fig. 81.

## LAN Failures and Split Group Actions

The PMU Server will not allow any cabling requests, destroy partition or actions on all partitions when the system is in a split state. The Partition Coordinator, when starting up after a new group is formed, determines if some of his members (EV7s and IO7s) are no longer available but were in the previous group. The requests that he accepts depend on the following Boolean conditions:

1. The Partition is running SRM or OS
2. Complete set of EV7s for the partition can be found in the group and perhaps only IO7s are incomplete
3. Whether the partition coordinator is contained in a minority or a majority group.

The table of Fig. 82 indicates the conditions when the Partition Coordinator can accept/reject the partition related commands. (T=True; X=Don't care)

## Discovery/Membership/Groups

MBMs and PBMs need to know about each other's existence. Some activities are restricted on degraded systems.

## Reconciliation concepts

When MBMs and PBMs form a group, they exchange their copies of the replicated databases. These databases may be inconsistent. The methods for arriving at a common replicated database among all the MBMs/PBMs are called reconciliation methods.

### Database reconciliation methods.

#### Case

Initial Power-on into a majority group  $\text{MaxPrevMajorityGroup} = (0,0)$

#### Merge Algorithm

All members load their database from NVRAM

The Leader obtains all database copies and compares, looking for a majority of copies that match. If there is no consensus, then the default values are used. The result is sent to all members via the first train of this group.

#### Case

Majority Group Formation with  $\text{MaxPrevMajorityGroup} \neq (0,0)$

#### Merge Algorithm

Retrieve db, PrevMajorityGroup, and the update list from each member whose PrevMajorityGroup matches the MaxPrevMajorityGroup. Use the copy with the longest update list as the initial state. The result is sent to all members via the first train of this group.

#### Case

Minority Group Formation with  $\text{MaxPrevMajorityGroup} \neq (0,0)$

#### Merge Algorithm

The db and updates are held. The train does not run. Access to the db (read or write) is blocked.

Important properties of the two phase train protocol:

1. The initiator of an update can only consider the update successful when it sees the train come around the first time. That means that all members have a copy of the unstable update. Thus, any resulting majority group will commit that update upon reconfiguring.
2. Updates initiated by members of the majority group are applied at most once by all members.

## Database Structures

The database information is shared by all MBM/PBM participants and contains a set of structures required to indicate the current operational state of a subset or the entire Multiprocessor computer system. Some database entries are maintained in a volatile form, others non-volatile and most in both forms. The reliable train messages are often used to communicate changes to the database copies of all members. The table of Fig. 83 indicates operations that cause modifications to databases maintained in RAM or NVRAM and the affected structures.

Turning to Fig. 84A and Fig. 84B, tables showing the structure of the database 84,000 is shown.

Components of the data base replicated system wide are maintained in flash memory, as shown in tables 84,002, 84,004 and 84,006, and 84,008. Table 84,002 gives the "Non-Volatile Partition Database". In table 84,002 the column "N/S" (or North - South) and the columns "E/W" (East - West) contain the mesh coordinates of the EV7 processors. These coordinates are assigned to the EV7 processors by the thumbwheels on each rack. These N/S and E/W values are the coordinates of the EV7 processors, and are the "key" to the database. The column "Hard" gives the partition number to which the EV7 processor is assigned by the human operator, and the value ranges from 0 to 255. The column "Sub" gives the subpartition to which the EV7 processor is assigned by the human operator. Table 84,002 is "Dense", or densely populated in that only entries for user defined EV7 processors are in the table.

Table 84,004 gives data on the microprocessor set. The column labeled "Rack" gives the rack number as read from thumbwheel settings, with the range of 0-32, as 32 racks may be assembled into a multiprocessor system. Each rack may hold from 0 to 8 EV7 processors. The

column labeled "Box" gives the box number of the IO boxes in the rack, and are set on the thumbwheels.

Table 84,006 gives the SRM environmental variables "SRM ENV VARs" arranged according to subpartitions. Collection of the SRM environmental variable information is given in Fig. 67 and the accompanying disclosure.

Table 84,008 gives the delta time "Delta Time" arranged according to subpartition. Delta time is described with reference to Fig. 63 showing a process for obtaining a new "base time", and in the accompanying discussion. Delta time is further discussed in the section entitled "Virtual TOY Clock".

Table 84,020 gives the database entries maintained in RAM memory by the MBMs and PBMs. Columns labeled N/S and E/W give the EV7 processor identifications, as also given in table 84,002. The columns identify the EV7 processor described by the horizontal entry in the table. The column headed "Hard" gives the hard partition to which the EV7 is assigned. The column "Sub" gives the subpartition to which the EV7 is assigned. The column labeled "Primary" states whether or not the EV7 is assigned as "primary" processor in its partition or subpartition (Y) or is not assigned as primary (N).

Cable connections of the EV7 are given in the entries headed "Cable Connections". The entries under "N" give the (x,y) coordinates of the EV7 connected to the processors North port. The column headed "S" gives the (x,y) coordinates of the EV7 connected to the processors South port. The column headed "E" gives the coordinates (x,y) of the EV7 connected to the processors East port. The column headed "W" gives the EV7 coordinates (x,y) connected to the processors West port. The column headed "I" gives the thumbwheel setting of the processor. The column headed "PID" gives the identification number of the processor, and ranges from 0-255. A cable connection value of "-1" indicates that the cable is disconnected. This table is dense, meaning that only the user defined EV7 processors are entered into the table.

Table 84,030 gives routing information to each processor for one EV7 processor in one hard partition. Table 84,030 gives the route through RBOXs for each processor in the multi-processor computer system to reach every other processor by a message. A copy of the data shown in table 84,030 is generated for each EV7 in a hard partition, and contains an entry to all other EV7s in the hard partition. Table 84,030 is divided into a section for each EV7 processor,

as indicated by (x,y) which give the processor's coordinates within the mesh. The processors identification (PID) is given as the Logical PID. The Logical PID identifies the processor entry, and gives the correspondence between the (x,y) entry giving the location of a processor within the mesh and the Logical PID. The processor identification is then followed by a route to each EV7 processor which it can reach in its partition. The route to each processor is made up of components such as: #1 RBOX\_ROUTE, #2 RBOX\_ROUTE, etc.

Each MBM handles up to eight (8) EV7 processors, and each EV7 may be in a different partition. Accordingly, each MBM may be required to maintain routing information for eight (8) partitions. The routing information maintained in table 84,030 is maintained in MBM RAM memory, and is regarded as temporary, or short lived information because a route between EV7 processors can change as a RBOX fails, a cable is disconnected, a new rack is hot swapped into the computer system, etc.

Table 84,040 maintains the status of each EV7 processor controlled by the MBM in question. Each MBM unit controls up to eight (8) EV7 processors, and so there may be up to eight entries in table 84,040. Table 84,040 is maintained by each MBM, and in an exemplary embodiment of the invention, is replicated throughout the MBMs of the partitions in which the EV7 processors are assigned.

Table 84,050 maintains information on "Base Time" of the multiprocessor computer system. The time base is the date, including day, month, year, hour, minutes, and seconds. This date information is maintained much like the time in any industry standard time of day chip.

## The PBM

The PBM hardware shown in Fig. 85 and software is much like that of the MBM. It does not have the quad uart. The PBM firmware block diagram shown in Fig. 86 is similar to the MBM firmware.

The firmware overview of Fig. 86 is the same as the MBM figure, except for the addition of an FPGA. Also the dark hashed areas indicate MBM functionality that does not exist on the PBM. The PBM has the same data bases as the MBM, it has the same group participation, the same capabilities to the partition coordinator, DHCP server, etc...

The services that an MBM provides to a CMM are listed as follows:

- Proxy Forwarding
- CMM presence watchdog
- TFTP Server
- Error repository & distribution
- Power hierarchy
- Time Services
- Virtual Console Terminal
- SRM Env. Var. Repository

The PBM does not have to do proxy forwarding, CMM watchdog, TFTP Server a PPP connection, store errors for a CMM, provide a time, or a virtual console, and there is no power hierarchy. The PBM is involved with SRM environment variables, since those variables are stored in the replicated non-volatile data base.

Although the PBM has an OCP, the OCP capability circle is reduced. The OCP functionality of a PBM refers to a power on/off switch and some LEDs.

#### FPGA Setup

The FPGA must be written with FPGA firmware, by the PBM. The PBM cannot load the firmware while the IO7 is active. The PBM only loads the FPGA firmware as part of a POWER ON command. Loading a test version of the FPGA firmware is not allowed. All IO7s must run the same FPGA code.

#### Set PCI Slot Information

The PBM does not have access to the IO busses. In order to help the PMU display information about the IO devices, the operating system probes the IO busses and provides this information to the SM subsystem. The operating system issues an SM protocol packet, SET PCI SLOT INFORMATION, with IO bus information. The packet destination is the PBM where the IO devices reside. The data is stored in the PBMs ram and can be retrieved by the PMU with a GET PCI SLOT INFORMATION.

## Error Messages and Actions

The table of Fig. 87 gives an embodiment of the messages and actions to be taken in reporting failures. It includes a series of fields that define the actions and additional localizing data that is stored in the error log.

Private LAN message routes and message formats are given in Figs. 88-110.

As shown in Fig. 88, network packets are passed between the EV7 and CMM via data buffer rings stored in CMM RAM. This allows a CMM to obtain requests from an application running on an EV7 and transfer responses to an EV7 from the MBM. The operating system can use this method directly. SM LAN packets transferred via these command rings do not contain IP/UDP headers. The CMM firmware will add/remove IP headers if these packets need to traverse the LAN.

As illustrated in Fig. 89, the first serial port on the CMM is used to communicate with the point to point serial port on the MBM at approximately 384,000 baud. The PPP protocol (RFC 1661) is used to transfer IP packets between the two microprocessors. The CMM implements a minimal PPP/IP/UDP stack in order to accomplish this.

As illustrated in Fig. 90, the CMM and EV7s all get IP address assignments according to the addressing scheme described in paragraph 3.5 *IP Addressing*. The MBM must respond to all requests on the LAN made to any of its own CMMs and EV7s and forward any requests to the appropriate CMM. This is done via two methods. The CMMs appear visible on the LAN via proxy ARP (RFCs 925 & 1027). EV7s are visible from the LAN via the MBM aliasing the EV7's IP address on its own Ethernet interface. The vxWorks *ifAddrAdd* function is used to install an IP address for each destination EV7 on the MBMs network interface.

Requests made from the EV7 side for another MBM will be forwarded directly to the member in question. The CMM forwards any SM LAN packets that it receives from the EV7, (that are not destined for itself), to the MBM.

As illustrated in Fig. 91, the CMM participates in the forwarding of requests from the EV7 to the MBM who handles the processing of the request and passes the response back to the CMM for forwarding to the EV7. For example, an APP running under UNIX may want to probe for the temperature reading in its IO box. The flow is depicted in Fig. 91.

As illustrated in Fig. 92, in general, the FPGA registers used for EV7 to CMM serial communication consist of: eight (8) single byte data registers for both transmit and receive on all four (4) virtual ports; and, two (2) control status registers - one (1) for each of the two (2) EV7s with bits indicating which data registers are full and if interrupts are enabled. The control status register also contains bits to indicate if interrupts occurred to the CMM from the EV7 and to request an EV7 interrupt from the CMM.

As is illustrated in Fig. 93, The EV7 that is designated as the primary for the partition uses the virtual COM ports to forward console characters to and from the MBM. This allows for a serial console pass through to a remote telnet connection running at the MBM. The PMU running a telnet session, on the SM LAN, connects to telnet session server on the MBM. There can be a telnet session for each COMx port that is being accessed.

The user types a character at the PMU telnet window. The character travels over the SM LAN to the telnet session on the MBM. This telnet session server envelops the character in a PUT\_CHAR protocol packet and passes it to the serial protocol driver task. The serial driver passes the packet to the CMM serial driver. The CMM serial driver strips off the envelope and writes the character to the COMx port.

As is illustrated in Fig. 94 and Fig. 95, messages passing through the Server Management Subsystem have a common header format identifying a command request and its associated response. Packets are encapsulated using UDP. For the most part any request expects to have a corresponding response; however, some Command Requests may require use of a broadcast, which only applies to those members that are found on the Server Management Subsystem's private LAN.

The common header on all messages contains the originator and destination of the request/response, a sequence number, a command code and the actual data contents. In the case of a response the top order bit of the command code is set and a positive or negative error reply immediately follows this command code. The length field will not exceed a predetermined maximum size and for cases where it is necessary to move data beyond that length a series of related command codes affords such movement at the application level.

Turning now to Fig. 94, the common Request message format is shown. Turning now to Fig. 95, the common Response message format is shown. Terms used in the message formats have the following meanings.

**Originator** - The IP address, associated with the requester of the data. If the requester is an OS application, the MBM places the IP address for the EV7 from which the request arrived. MBMs, PBMs and PMUs have pre-assigned IP addresses for the LAN. CMMs and EV7s have fixed IP addresses even though they are not on the LAN, but the MBM that they are attached to shall respond and forward requests for them.

**Destination** – The IP address of the target expected to respond to the command. In the case that a broadcast is used the destination field contains the special naming convention: broadcast 255.255.255.255. Broadcast destination addresses are only meant for those entities on the Server Management Subsystem private LAN. Another special naming convention is used when an OS application wants to address the closest MBM, the address 0.0.0.1 is used.

**Identifier** – a reserved space used by the application to match requests with responses.

**Command Code** - consists of a request response bit, a group identifier and a request identifier. The bit breakdown of this byte is defined as follows: bit 15, 0-request, 1-response; bit 14-8, message group identifier; and, bit 7-0. request identifier.

Messages on the LAN are sent via UDP packets using one of the sockets listed below: UDP port No. 710, Group messages; UDP port No. 711, SM LAN Command/Response Traffic; UDP port No. 712, Atomic Update Messages; and, UDP port No. 323-335, Virtual Console traffic, assigned by GET TELNET IP ADDRESS PORT request.

**Result Code** - 0 – no error; else refer to list of errors given in Fig. 87.

**Data Content** - The values depend on the command in use.

Group Messages are the messages used to form or maintain the group communication protocol. These messages include those that are broadcast to all microprocessors directly connected to the LAN.

SM LAN Command/Response Traffic, is any of the commands or responses described in the succeeding chapters. These are the “action” commands of the SM LAN.

Atomic Update Messages are those that are used to perform guaranteed communication and/or replicated database updates through the train protocol.

Virtual Console traffic is the communication that implements the MBM to CMM link.

Referring now to Fig. 96, the train header is used on messages that are performing operations and/or updates on all group members in a guaranteed, ordered fashion. The general algorithm for the train protocol is described. The first train message is “injected” onto the LAN by the leader upon successful group formation. The train message begins with a Command Code of “empty”. When a member wishes to initiate a command or update for all members, it waits to receive the empty train. Before sending the train message to the next member, it changes the command code to “full” and appends the command packet. When the message is received by the initiator a second time, it changes the command code to “empty” and removes the command packet from the message.

Format:

Originator IP Address – should be 0 – special case for a train message.

Destination Node – should be ffffffh – special case for a train message.

Identifier – should be the GroupID – allows matching to detect a left over train message from a previous group.

Command Code - xxxx = train empty , yyyy = train full. Indicates if there is a valid SM LAN message appended.

SM LAN message – optional appended message to be executed by all nodes.

Train messages are MBM or PBM member to member commands / updates only. They can not be initiated or received by the PMU, CMMs, or EV7s. They are internal to the SM LAN and any usage by external utilities is prohibited.

Any SM LAN message can be sent via the train. However, not all messages are sent via train mechanism.

SM LAN messages sent via train messages are not replied to. Successful travel around the ring of members twice implies success of the command / update. Messages that are not guaranteed to succeed should not be sent via this mechanism, but should be sent via a unicast method from the leader (or other originating node).

Referring now to Fig. 97-110, message type groups are given, along with the name of the messages in each group.

Fig. 97 gives the messages in the LAN Formation Group. The LAN formation group of messages is used to implement the group communications protocol. This set of messages is only implemented by the MBM and PBM microprocessors.

Fig. 98 gives the Reliable message group. These messages define the messages for the train protocol.

Fig. 99 gives the system discovery group of messages. These messages derive the LAN components configuration and current state.

Fig. 100 gives the Partition Control group of messages. This group of messages has commands that effect changes in the state or configuration of partitions, including migration of EV7 processors, power on/off, and halt CPU.

Fig. 101 gives the EV7 Setup group of messages. This group of messages consists of commands to manipulate the EV7, aid in the manual configuration of IO7 cabling, and the EV7 CPU assignment for the IO7 subsystems.

Fig. 102 gives the Cable Test group of messages. The Multiprocessor computer system contains test signals in each cable linking pairs of EV7 routing axes around the external 8P connections. The 6 IP cables link EV7s for routing purposes and the 8 IO cables link the single EV7 to IO7 risers on PBMs. The signals light and extinguish LED pairs and may be asserted and/or read by MBMs or PBMs at each end of the cable. This command is directed at an MBM or PBM to modify the state of the cable test signals controlled by that entity.

Fig. 103 gives the Virtual Console group of messages. When the application requests a virtual console window, a telnet session is established between the requester and the MBM associated with the primary EV7 for the partition. Until the session is closed, all keystrokes and console output data are passed between MBM, CMM and the primary EV7 of that partition and handed over to the telnet session. Each primary EV7 has 2 virtual COM Port connections. Obtain the proper session by making a GET TELNET IP ADDRESS/PORt request. Opening a telnet session should cause the last 2k of buffered output from the virtual console to be displayed.

Fig. 104 gives the Firmware Load and Upgrade group of messages. Firmware for SROM, XSROM, FPGA, Memory Test, CMM, SRM, MBM and PBM are stored in the CMM, MBM or PBM Flash area. These commands facilitate the upgrading, loading and testing of Firmware for the Multiprocessor computer system Server Management subsystem.

The Trivial File Transfer Protocol (TFTP) is used to transfer flash image “files” throughout the Server Management LAN. These files are actually segments of the flash ROMs on each of the MBM, PBM, and CMM modules. The MBMs implement a simple TFTP server to provide images that are requested by the CMM. The PMU contains a TFTP server to provide Fail Safe Loader (FSL) images to MBMs, PBMs or CMMs that request them.

Fig. 105 gives the Environmental Retrieval group of messages. Application requests are made to the Server Management Subsystem to obtain dynamically changing data and states.

Fig. 106 gives the Field Replaceable Unit (FRU) group of messages. Serialization and other manufacturing data are stored in EEROMs on several Server Management Subsystem controlled components. Application programs use these requests to retrieve and set field values in these EEROMs.

Fig. 107 gives the error logging group of messages. Error logging takes on several forms. Some error information is maintained in the EEROM of the device itself (power supplies and

RIMMs); some is reported on the OCP and some errors are logged to the NVRAM of the MBM, PBM or CMM that have the offending error. Errors can also be reported to the PMU and OS Application in the form of an alert.

Fig. 108 gives OS Watch Dog Timer group of messages. An OS Application can make a request to start a Keep Alive Watchdog. The MBM receiving this request expects to receive a keep alive message within the application indicated time frame or take the action(s) indicated in the start message. Before an OS shuts down, the Stop Watchdog Timer request is made.

Fig. 109 gives Date/Time group of messages. The Server Management LAN provides a time service that is used to provide the Alpha SRM BB\_WATCH functionality. This takes the form of a Base Time that is shared by all microprocessors and maintained in a battery backed up watch chip, and a Delta Time that is maintained on a partition basis. The commands in this section are used to get and set the Base and Delta Time values.

Fig. 110 gives the Miscellaneous group of messages. These commands provide various utility functions.

It is to be understood that the above described embodiments are simply illustrative of the principles of the invention. Various other modifications and changes may be made by those skilled in the art which embody the principles of the invention and fall within the spirit and scope thereof.

What is claimed is:

## CLAIMS

1       1. A multiprocessor computer system having a plurality of processors interconnected so that  
2       they can share memory, comprising:

3               a plurality of links, each link of said plurality of links connecting a processor to another  
4       processor;

5               a router box (RBOX) associated with each processor of said plurality of processors, said  
6       RBOX arranged to forward a message received on an input link of said plurality of links from a  
7       source processor to an outgoing link of said plurality of links toward a destination processor in  
8       response to data carried in said message;

9               a plurality of microprocessors, each of said microprocessors having a microprocessor  
10      memory associated therewith, a selected microprocessor of said plurality of microprocessors as-  
11      sociated with at least one processor of said plurality of processors, said plurality of microproces-  
12      sors arranged to control said plurality of processors, said control including applying electric  
13      power to a selected processor and removing electric power from said selected processor;

14               a data structure stored in microprocessor memory, said data structure storing a represen-  
15      tation of the links connecting said processors and storing routes used by said RBOX in routing  
16      messages along said links, a copy of said data structure stored in microprocessor memory of each  
17      of said microprocessors; and,

18               a process to update said data structure in each said microprocessor memory in the event  
19      that a change occurs in a status of a component of said multiprocessor computer system.

1       2. The apparatus of claim 1 further comprising:

2               a process executing in said microprocessors for directing said microprocessors in forma-  
3      tion of a partition of said processors, said partition having a selected number of said processors  
4      as members, said members capable of reading and writing a common memory within said parti-  
5      tion, and other non-member processors excluded from reading and writing said common mem-  
6      ory; and,

7        a second data structure for storing a representation of said partitions, and storing routes  
8 through said links for transfer of messages between processors within a partition but not between  
9 processors of different partitions.

1        3. The apparatus of claim 2 further comprising:

2            a management computer communicating through said local area network with said mi-  
3 croprocessors, said management computer having an input device such as a keyboard and mouse  
4 for entering commands to said plurality of microprocessors to modify said data base in order to  
5 establish the processors belonging to a partition, said multiprocessor computer system supporting  
6 a plurality of said partitions;

7            a process responsive to said data base, said process executing in said microprocessors to  
8 establish that a processor in a partition can read and write memory associated with other proces-  
9 sors of said partition, but cannot read and write memory associated with processors which are not  
10 members of said partition.

1        4. The apparatus of claim 1 further comprising:

2            an input/output (IO) subsystem associated with selected processors of said plurality of  
3 processors;

4            an IO microprocessor associated with each said IO subsystem, said IO microprocessor  
5 communicating with said microprocessors through said local area network, each said IO micro-  
6 processor having an IO microprocessor memory holding a copy of said database;

7            a process executing in said IO microprocessor, said process responsive to said database,  
8 said process to apply power or remove power to said IO subsystem.

1        5. The apparatus of claim 1 further comprising:

2            a boot-up process executing in microprocessors of said plurality of microprocessors, said  
3 boot-up process to start execution of said processors of said multiprocessor system, and said  
4 multiprocessor system capable of continuing operation, after start of execution of said proces-  
5 sors, to permit removal of a microprocessor from the multiprocessor system without interrupting  
6 execution of said processors.

1       6. A method for operating a multiprocessor computer system, comprising:  
2              connecting a plurality of processors so that they can share memory, a processor of said  
3              plurality of processors connected to another processor by at least one link of a plurality of links;  
4              associating a router box (RBOX) with each processor of said plurality of processors, said  
5              RBOX arranged to forward a message received on an input link of said plurality of links from a  
6              source processor to an outgoing link of said plurality of links toward a destination processor in  
7              response to data carried in said message;  
8              interconnecting a plurality of microprocessors, each of said microprocessors having a mi-  
9              croprocessor memory associated therewith, a selected microprocessor of said plurality of micro-  
10             processors associated with at least one processor of said plurality of processors, said plurality of  
11             microprocessors arranged to control said plurality of processors, said control including applying  
12             electric power to a selected processor and removing electric power from said selected processor;  
13              storing a data structure in microprocessor memory, said data structure storing a repre-  
14              sentation of the links connecting said processors and storing routes used by said RBOX in rout-  
15              ing messages along said links, a copy of said data structure stored in microprocessor memory of  
16              each of said microprocessors; and,  
17              updating said data structure in each said microprocessor memory in the event that a  
18              change occurs in a status of a component of said multiprocessor computer system.

1       7. The method of claim 6, further comprising:  
2              executing a process in said microprocessors for directing said microprocessors in forma-  
3              tion of a partition of said processors, said partition having a selected number of said processors  
4              as members, said members capable of reading and writing a common memory within said parti-  
5              tion, and other non-member processors excluded from reading and writing said common mem-  
6              ory; and,  
7              storing in a second data structure a representation of said partitions, and storing routes in  
8              said second data structure through said links, said routes for transfer of messages between proc-  
9              essors within a partition but not between processors of different partitions.

1       8. The method of claim 6, further comprising:

2 establishing communication between a management computer and said microprocessors  
3 through said local area network, said management computer having an input device such as a  
4 keyboard and mouse for entering commands to said plurality of microprocessors to modify said  
5 data base in order to establish the processors belonging to a partition, said multiprocessor com-  
6 puter system supporting a plurality of said partitions;

7 establishing, in response to said data base, that a processor in a partition can read and  
8 write memory associated with other processors of said partition, but cannot read and write mem-  
9 ory associated with processors which are not members of said partition.

1 9. The method of claim 6, further comprising:

2 associating an input/output (IO) subsystem with selected processors of said plurality of  
3 processors;

4 associating an IO microprocessor with each said IO subsystem, said IO microprocessor  
5 communicating with said microprocessors through said local area network, each said IO micro-  
6 processor having an IO microprocessor memory holding a copy of said database;


7 executing a process in said IO microprocessor, said process responsive to said database,  
8 said process to apply power or remove power to said IO subsystem.

1 10. The method as in claim 6, further comprising:

2 executing a boot-up process in microprocessors of said plurality of microprocessors, said  
3 boot-up process to start execution of said processors of said multiprocessor system, and said  
4 multiprocessor system capable of continuing operation, after start of execution of said proces-  
5 sors, to permit removal of a microprocessor from the multiprocessor system without interrupting  
6 execution of said processors.

**ABSTRACT**

The invention is a control system using microprocessors which communicate through a Local Area Network (private LAN) to control operation of both processors and input and output subsystems (IO system) of a multiprocessor computer system. The processors each have memory associated therewith, and each processor has an IO system comprising a plurality of busses such as PCI busses, associated therewith. The processors are cabled together in a mesh arrangement so that messages can be transferred between any of the processors and delivered to memory associated with the destination processor, or delivered to an IO system associated with the destination processor, etc. The microprocessors are powered on when power is applied to the chassis of the multiprocessor system, and the microprocessors then control the processors of the multiprocessor system, including applying power to the processors, forming hard partitions containing selected processors, computing routes from a processor to a memory associated with any processor for read and write transactions, computing routes to IO subsystems associated with any processor of the hard partition, forming partition boundaries so that processors in one hard partition cannot read and write to memory or IO systems associated with processors in another hard partition, forming soft partitions of processors, controlling boot-up of operating systems executing on the processors of the multiprocessor computer system, removing power from a failed processor, providing power to a repaired processor, etc.



1  
FIG.




FIG. 2




FIG. 3




FIG. 4




FIG. 5

+ -




FIG. 6

© 1995-2000 by the Board of Education of the City of New York

-4-





FIG. 7

FIG. 8






FIG. 9




FIG. 10

<b>Command</b>	<b>Description</b>
connect	Connect a virtual console session to an SRM Firmware instance
disconnect	Disconnect a virtual console session from an SRM Firmware instance
power {off, on} <i>item</i>	Change the power state of an item (CPU, 8P unit, I/O Drawer)
halt <i>processor</i>	Issue a HALT to an EV7 processor (specified by processor ID)
reset [ <i>partition</i> ]	Issue a system reset or a reset to the specified partition
reset <i>processor</i>	Issue a Reset to an EV7 processor (specified by processor ID)
reset { CMM, MBM, PBM}	Issue a Reset to one of the Server Management processors
set manufacturing	Set serial number, other FRU data
set partition	Define the number of partitions, assign partitionable resources to each partition, define the partition permissions.
show config	Show the entire system configuration
show CPU { <i>processor</i> }	Show data on an EV7 CPU (specified by processor ID) or all CPUs
show LAN	Show the nodes on the Server Management LAN
show memory	Show information on memory configuration
show partitions	Show the defined partition data
show FRU	Show the FRU data for the system FRUs
show power	Show the thermal and voltage sensor data
show error	Show the non-volatile saved error state
clear error	Clear the non-volatile saved error state
update	Update system firmware
date	Set / show the server management time
examine / deposit	Display / modify memory
test <i>processor n</i>	Run the test identified by <i>n</i> on the specified <i>processor</i> ID
set test signal <i>processor n</i>	Assert cable test signal for port n (N, S, E, W, I/O) on the specified <i>processor</i> ID and light the cable LED.
clear test signal <i>processor n</i>	De-assert cable test signal for port n (N, S, E, W, I/O) on the specified <i>processor</i> ID and extinguish the cable LED.
check test signal <i>processor n</i>	Test cable test signal for port n (N, S, E, W, I/O) on the specified <i>processor</i> ID.

*External Server Management Commands*

Fig. 11

<i>Command</i>	<i>Description</i>
PutChar	Send a character to the operator display
GetChar	Get a character from the operator keyboard
SetTermInt	Set operator terminal interrupt setting
Hello	Announce the present of a Server Management member
Poll	Probe for the presence of a specific Server Management member
No-op	No operation, used for testing
SysError <i>data</i>	System error state information to be saved
FRUError <i>id</i> , <i>data</i>	Store FRU error <i>data</i> in the FRU specified by <i>id</i>

*Internal Server Management Commands*

Fig. 12

<i>Cell</i>	<i>Description</i>
Seconds	Second count, 0-59, binary format
Minutes	Minute count, 0-59, binary format
Hours	Hour count, 0-23, binary format
Day	Day of the month, 1-31, binary format
Month	Month of the year, 1-12, binary format
Year	Year, 0-99, binary format


*'BB\_WATCH Data*

Fig. 13

<i>Message</i>	<i>Description</i>
CSB_READ	Read data element from the PBM
CSB_WRITE	Write data elements to the PBM
CSB_POLL	Obtain PBM status

*CSB Messages*

Fig. 14



*Server Management Hardware Overview*

Fig. 15





Fig - 16




*Rack and Box Thumbwheel Scheme*

Fig. 17

Task	How Many	Where
Group Leader	1 per Marvel System <sup>1</sup>	Lowest MBM in group
PMU Server	1 per Marvel System <sup>1</sup>	Lowest MBM in group
Partition Coordinator	1 per hard partition; max 8 per MBM	MBM with lowest EV7 in hard partition
Telnet Server	2 per subpartition (COM0,COM1)	Grandparent MBM of primary EV7
DHCP	1 per Marvel System <sup>1</sup>	Lowest MBM in group

*MBM Task Attributes*

Fig. 18






Fig. 20


General Phase	Group Changes	Specific Steps	Notes
Independent Processing	Unaffected	POST H/W Poll Status and State	Handle MBM hot swap. Power that is on is left on, power that is off is left off.
Group Formation Majority / Minority Replicated Data Sync	Normal group processing, new group is formed	New Group is initiated. A group forms and a leader is selected. If there was not a previous majority group, then the replicated database is marked invalid. If the new group is a minority, mark the database read-only. If the new group is a majority, request the database and any pending updates from all members who were previously joined to the MaxPrevMajorityGroup. Apply the longest list of updates to the corresponding database copy and send the new initial replicated database to all members. All members mark the database as valid. If the new group is a majority and there was no previous majority group, then clear the powerup_complete flag. If the new group is a majority and the powerup_complete flag is set, then proceed to the Hardware Init phase. Else, proceed to the Operational phase.	Minority groups remain in this phase
Hardware Init	Return to forming a new group and re-run init	Poll CMMs to determine CPU module population details within each 8P backplane. This initializes the list of available resources. The Leader computes a routable configuration for each partition based upon the requested and available CPU resources. The leader decides to power up the partitions. It commands all MBMs, PBMs, and CMMs to power up. Upon completion, the leader obtains the results from each of these commands and adjusts the list of available resources accordingly. The partition coordinators start XSROM testing on all CPUs for memory. The PMU Server initiate IP cable testing between 8P backplanes and IO cable testing between 8P backplanes and I/O crates. The results are used to recalculate routing and assign I/O to partitions. The partition coordinators initiate XSROM tests for I/O and routing. The results are used to recalculate routing and adjust the list of I/O resources. The partition coordinators initiate remote memory testing between CPUs in the same partitions. The partition coordinators initiate interrupt testing between CPUs in the same partitions. The leader sets the powerup_complete flag to true and proceeds to the S/W Load phase.	
S/W Load (SRM + O/S)	New Groups are treated as hot-adds	The partition coordinators elect a primary EV7 in each partition. They command all secondaries to spin on RBOX_SCRATCH and initiate loading of the SRM firmware on the primary. The SRM firmware commands the secondary CPUs to join by writing RBOX_SCRATCH. The primary EV7 completes all I/O initialization. If the SRM auto_action environment variable is set to BOOT, the operating system boot is attempted on the partition.	
Operational	New Groups are treated as hot-adds	Server management requests from the primary EV7 in each partition are handled by the CMMs / MBMs / PBMs. If a new group has caused a change in the CPU, Memory, or I/O resources, notify the primary EV7 in each affected partition. If a new resource has been pre-allocated to a partition, then the partition coordinator takes the steps necessary to probe the IP or IO links to the new resource (CPU or I/O drawer).	

Marvel System Powerup Flow with Group relationship

Capability	via private LAN	via FPGA
Virtual console terminal access	Yes	No
Firmware updates	Yes	No
Load/Disable Test firmware	Yes	No
Live configuration change <sup>1</sup>	No	Yes
Writing SRM environment vars <sup>2</sup>	No	Yes
Unsolicited notification of alerts	Yes	No
Store PCI Slot information <sup>2</sup>	No	Yes
All others	Yes	Yes

*LAN vs FPGA PMU capability matrix*

Fig. 22



*PMU Server Block Diagram*

Fig. 23

Command Group	Commands Received	Class	PMU Server Handling Method
System Discovery	Get MBM / PBM CONFIGURATION	Forward	Commands are forwarded to the MBM or PBM addressed in the message header. Responses are forwarded back to the PMU.
	GET PCI SLOT INFO		
	GET PARTITION DATABASE	Direct	The PMU Server derives the response from his local copy of the Partition Database.
	GET OWN PARTITION NUMBER		
	GET SYSTEM TOPOLOGY	Direct	The PMU Server keeps track of the applications making this multi phase request until the last entity has been requested. The entity number in the request is used to index into a list composed of the combination of group members and partition database. The IP addresses, parent relationship and partition number is derived from these values found in NVRAM. If the group members or partition database changes before the last entity is requested, an error response is returned on the next request.
EV7 Setup	SET PCI SLOT INFO	Forward	The PMU Server must ensure that this request is coming from an EV7 and not from the PMU on the LAN. These packets are directed to the PBM associated with the slot.
	ALL COMMANDS OF GROUP	Forward	All commands in this group that contain a partition number are forwarded to the MBM running the appropriate partition coordinator. The exceptions are Read State of OCP switches, OCP Switch assignment and Power On/Off commands that are simply forwarded to their destination.
Cable Test Group	REQUEST EV7 START TEST	Forward	This command is forwarded to the destination EV7.
Virtual Console	GET CABLING CONFIGURATION	Direct	The PMU Server responds with the contents of the Cable Database that he composed during Initialization or was requested via Reconfigure Cabling.
	RETEST CABLE CONFIGURATION	Complex	The PMU Server re-initiates the test of all IP and IO cabling making use of the commands Get MBM IP Cabling and Get PBM IP Cabling.
	SET CABLE TEST SIGNAL	Complex	The PMU Server ensures that there are no other on-going cabling requests and forwards these commands to the PBM or MBM in the destination field and returns the response to the PMU.
	GET CABLE TEST SIGNAL		
GET TELNET IP ADDRESS/PORT	Direct	PMU Server determines the primary MBM's IP Address and the socket port.	
Firmware Load and Upgrade, Environmental Retrieval, FRU Data, Error Logging, Miscellaneous		Forward	All commands in these groups are forwarded to the CMM, PBM or MBM in the destination field and the response returned to the PMU.
Date/Time		Forward Direct	The PMU Server allows Base Time Gets and Sets from all PMUs but Delta Time Sets and Gets can only come from Partition Primary EV7s. The requests are forwarded to the MBM being addressed.
Miscellaneous	GET/SET KNOB READ/WRITE BLOCK DATA	Forward	These requests are forwarded to the destination for processing.


PMU Server Received Command Handling

Fig. 24

Command Group	Commands Received	PMU Server Handling Method
System Discovery	Get MBM/PBM CONFIGURATION	
	GET PARTITION DATABASE	
Cable Testing	SEND CABLE ID RECEIVE CABLE ID GET MBM IP CABLING GET PBM IP CABLING	These commands are issued in response to a RETEST CABLE CONFIGURATION request. The process is discussed in section Error! Reference source not found., Error! Reference source not found.
Error Logging Group	ERROR REPORTING	The PMU Server knows the IP address of client PMUs and distributes the alerts to each PMU.
Miscellaneous	DISTRIBUTE DHCP LEASE DATA	The DHCP server runs on the PMU Server and keeps track of the DHCP clients. For failover purposes, this data is replicated on all nodes. See section Error! Reference source not found. Error! Bookmark not defined..

*PMU Server Originating Commands*

Fig. 25



*Fig. 26*

PMU	MBM PMU Server DHCP Server	CMM	EV70	EV71	PBM	The PMU's PC connects to the LAN and requests an IP address
PMU	MBM PMU Server DHCP Server	CMM	EV70	EV71	PBM	The MBM is running a DHCP server and provide it a lease on a DHCP address.
PMU	MBM PMU Server DHCP Server	CMM	EV70	EV71	PBM	The PMU application when initializing issues the GET SYSTEM TOPOLOGY command starting at 0. The command is issued to the PMU Server which is a predetermined address
PMU	MBM PMU Server DHCP Server	CMM	EV70	EV71	PBM	Entity 0, the data for the MBM is returned.
PMU	MBM PMU Server DHCP Server	CMM	EV70	EV71	PBM	PMU sends a GET MBM CONFIGURATION to the MBM and picks up information, like the RIMM population.
PMU	MBM PMU Server DHCP Server	CMM	EV70	EV71	PBM	The PMU Server task forwards the command to the MBM SM protocol servicing task. It returns information on the CPUs that is has gathered previously.
PMU	MBM PMU Server DHCP Server	CMM	EV70	EV71	PBM	The PMU issues the GET SYSTEM TOPOLOGY for entity 1
PMU	MBM PMU Server DHCP Server	CMM	EV70	EV71	PBM	Entity 1, the CMM, is returned.

*Show Config Flow Diagram (Part 1)*

*Fig. 27*

	MBM PMU Server DHCP Server	CMM	EV70	EV71	PBM	The PMU issues the <b>GET SYSTEM TOPOLOGY</b> for entity 2
	MBM PMU Server DHCP Server	CMM	EV70	EV71	PBM	Entity 2, the data for the EV7-0 is returned.
	MBM PMU Server DHCP Server	CMM	EV70	EV71	PBM	The PMU issues the <b>GET SYSTEM TOPOLOGY</b> for entity 3
	MBM PMU Server DHCP Server	CMM	EV70	EV71	PBM	Entity 3, the data for the EV7-1 is returned.
	MBM PMU Server DHCP Server	CMM	EV70	EV71	PBM	The PMU issues the <b>GET SYSTEM TOPOLOGY</b> for entity 4
	MBM PMU Server DHCP Server	CMM	EV70	EV71	PBM	Entity 4, the data for the PBM is returned.
	MBM PMU Server DHCP Server	CMM	EV70	EV71	PBM	The PMU issues a <b>GET MBM IP CABLING</b> to the PMU Server
	MBM PMU Server DHCP Server	CMM	EV70	EV71	PBM	IP Cabling data is returned

Show Configuration Flow Diagram (Part 2 )

Fig. 28

 PMU	MBM PMU Server DHCP Server	CMM	EV70	EV71	PBM	GET PARTITION DATABASE
 PMU	MBM PMU Server DHCP Server	CMM	EV70	EV71	PBM	PMU Server delivers the database
 PMU	MBM PMU Server DHCP Server	CMM	EV70	EV71	PBM	The PMU uses the information from the REQUEST COMPLETE LAN TOPOLOGY, and issues commands with IP addresses. It issues GET VOLTAGE READINGS to the MBM
 PMU	MBM PMU Server DHCP Server	CMM	EV70	EV71	PBM	The PMU Server passes the command to the destination, the MBM. The MBM responds with the voltage readings.
 PMU	MBM PMU Server DHCP Server	CMM	EV70	EV71	PBM	The PMU uses the information from the REQUEST COMPLETE LAN TOPOLOGY, and issues commands with IP addresses. It issues GET VOLTAGE READINGS to the PBM
 PMU	MBM PMU Server DHCP Server	CMM	EV70	EV71	PBM	The PMU Server passes the command to the destination, the PBM. The PBM responds with the voltage readings.
 PMU	MBM PMU Server DHCP Server	CMM	EV70	EV71	PBM	The PMU uses the information from the REQUEST COMPLETE LAN TOPOLOGY, and issues commands with IP addresses. It issues GET VOLTAGE READINGS to the CMM
 PMU	MBM PMU Server DHCP Server	CMM	EV70	EV71	PBM	The PMU Server passes the command to the destination, the CMM. The CMM responds with the voltage readings.
 PMU	MBM PMU Server DHCP Server	CMM	EV70	EV71	PBM	Steps are repeated for GET FAN RPM SPEED, GET TEMPERATURE READING, GET POWER SUPPLY, GET EEPROM DATA
 PMU	MBM PMU Server DHCP Server	CMM	EV70	EV71	PBM	The PMU Server queries the PBM for information on the IO drawers with the GET PBM CONFIGURATION command
 PMU	MBM PMU Server DHCP Server	CMM	EV70	EV71	PBM	The PBM supplies information on each IO7 Riser in the IO Drawer
 PMU	MBM PMU Server DHCP Server	CMM	EV70	EV71	PBM	GET PCI SLOT INFO
 PMU	MBM PMU Server DHCP Server	CMM	EV70	EV71	PBM	The PBM may have stored information on the PCI configuration if it was stored by SRM console in the PBM ram.
 PMU	MBM PMU Server DHCP Server	CMM	EV70	EV71	PBM	The process is complete.

Fig. 29 Show Configuration Flow Diagram ( Part 3 )

Configuration of Marvel at: 5/22/00 08:30:05

MBM Rack 3-01 Part No = A1234567, Revision = 32, Serial No = 1233456789 MFGDate 040101 3 Errors

CMM 0 Part No = C1135-66, Revision = 01, Serial No = 11335577

FV70 - Partition 8/0; GTL Voltage = 1.5, 5v = 4.9, Temp = 85F; Part No = E7-435, Revision = 5, Serial No = 9876543, MFGDate = 101100

FV71 - Partition 8/0; GTL Voltage = 1.0, 5v = 5.0, Temp = 45F; Part No = E7-435, Revision = 2, Serial No = 9876530, MFGDate = 092100

RIMM0 - 256MB; Part No = RA03-256, Revision = 1, Serial No = 000123456, MFGDate = 070199

RIMM1 - 256MB; Part No = RA03-256, Revision = 1, Serial No = 000123457, MFGDate = 070199

48V Power Supplies: 1=Operational, 2=Missing, 3=Operational, 4=Off

FAN1 = 340RPM(minimal=200),

PBM Rack 6-02 Part No = B1234567, Revision = 32, Serial No = 1233456790 MFGDate 040100 0 Errors

100W Power Supply 1=Operational, 5v = 5.0, 5v = 5.1, 12v = 11.99, 3.3v = 3.1, Temp = 98F

FAN1 = 280rpm(minimal=100), FAN2 = 330rpm(minimal=100); Part No = PA003, Revision = 01, Serial No = SS2358


(0) Drawer - Partition 8; Part No = CP000IO1, Revision = 07, Serial No = IO2345

PCI slot 3 - class = 01, subclass = 04, deviceid = 0701, vendorid = 0809, mfrin = A, irq = 5

PCI slot 3 - class = 02, subclass = 03, deviceid = 0e00, vendorid = a312, mfrin = D, irq = 9

Show Config Sample Output

Fig. 30



*Partition Request Sources  
Fig-3/*

Request Received	Class	Partition States	Handling Methods	Reference Flow
Move EV7s to Partition Remove EV7s from Partition	Complex	Not Running OS	Check the proper cabling of a newly formed hard partition or additions to an existing partition. Determine memory and routability. Distribute the requests on the TRAIN to update volatile Databases.	
Assign Memory or IO to Sub Partition	Trained			
Set Partition State Attributes	Trained			
Switch Primary EV7	Complex	Partition Running OS	Distribute the requests on the TRAIN to update volatile Database.	
Assign Sub Partitions	Trained			
Assign Memory or IO to Sub-partition	Trained			
Add EV7 or Delete EV7 to/from Running Partition	Complex	Partition Running OS	Cable testing is performed to determine proper connections ; A Quiesce request is made to all EV7s in the partition; new Rconfig/Cconfig requests are sent to all EV7s and if acceptable Continue is sent to the EV7s. If an error is found in routing new Rconfig/Cconfig computations are attempted and sent to the affected EV7s.	Figure 16
Save Partition Assignment	Trained	Don't Care	Send the save command on the train so that all copies of volatile database for the partition can be copied to non-volatile.	Figure 9
Reset Partition	Complex	Partition Reset in Progress or Partition running OS	Pulse Reset on all EV7s in partition, Reload SROM, XSROM on all EV7s, reconfigure Routing, send Rconfig/Cconfig to all EV7s and if no errors, Load SRM to primary.	Figure 11
		Partition Powered Off	Return Error	
Power Off Partition	Complex	Partition Reset in Progress or Partition running OS	If Dual EV7 in partition or EV7 in free pool, send Power Off EV7s to each such CMM in partition. Power off PCI Drawer on PBMs that have all IO7 risers in partition. If power is off for all CMMs in a cabinet of MBMs, power off the power supplies in the cabinet.	
		Partition Powered Off or request made to sub partition	Return Error	
Power On Partition	Complex	Partition Reset in Progress or Partition running OS	Return Error	
		Partition Powered Off	Turn Power on to any Power Supplies that are off and under control of Cabinets that have EV7s assigned to our partition. Power up PCI drawers for PBMs that have IO7s in our partition. Power up Dual EV7s on CMMs that have EV7s in our partition. Continue like the reset process.	
Halt Partition	Forwarded	Partition Reset in Progress or Partition running OS	Send Halt On to primary Ev7 of partition.	
		Partition in Halt State or Partition Power off	Return Error	
Disable Halt Partition	Forwarded	Partition Reset in Progress or Partition running OS	Return Error	
		Partition in Halt State or Partition Power off	Send Halt Off to primary Ev7 of partition	
Store Environment Variables	Trained	Partition running OS	This causes a distribution of the SRM environment variables to all peers	Figure 34
Get Environment Variables	Direct	Partition running OS	Return local copy of all SRM environment variables	

Partition Coordinator Handling of Requests

Fig. 32

Commands Issued	Partition States	Handling Methods	Reference Flow
Train Full	Any		
EV7 Quiesce	Partition Reset in Progress	CMM facilitates	Figure 17, Figure 20
Config Rbox/Cbox	Partition Reset in Progress	CMM facilitates	Figure 17, Figure 20
EV7 Reset	Any	CMM facilitates	Figure 15
EV7 Start Test	Partition Reset in Progress	CMM facilitates	Figure 16, Figure 18
EV7 Halt	Any	CMM facilitates	
Load Image	Partition Reset in Progress	CMM facilitates	Figure 15, Figure 20
Set Partition State	Partition Reset in Progress	CMM facilitates	Figure 15, Figure 16, Figure 17, Figure 18, Figure 20, Figure 21
Power On/Off	Any	CMM, MBM, and PBM facilitates	

*Partition Coordinator Commands Issued*

**Fig. 33**





Fig. 34



*Inputs and Outputs of Router algorithm*

Fig. 35

### *Routing Glossary*

Primary dimension: one of EAST-WEST or NORTH-SOUTH. This choice is the same for all EV7s.

Secondary dimension: the other way.


Dimension-order routing: the shortest path connecting two nodes which proceeds first along the primary dimension and then along the secondary.

Adaptive routing: the collection of paths advancing node-to-node in the same primary and secondary directions as the dimension-order routing. At each intermediate node it must be possible to advance in either direction until the dimension coordinate in a direction matches that of the destination.


Initial hop: a routing option which allows a hop from the source node in any direction to an adjacent node. This option allows some connection of nodes in imperfect meshes.

SRC routing: another deadlock-free routing method in which travel proceeds first along the secondary dimension.


Fig. 36 Routing Glossary




Partition Coordination  
Fog- 37




At initial poweron, the unconfigured Marvel system is configured to have all resources in a single subpartition. The global free pool (HP 255) is empty.



The user creates hard partition 1 (HP1) with subpartition 0 (SP 0). The partition free pool (SP 255) is created automatically.








The user deletes several EV7s from the hard partition (HP 0) and they migrate to the global free pool (HP 255).



The user can move the EV7s into the new partition, (HP 1,SP 0). By default the IO7s are assigned to the same partition as the EV7 to which it belongs.

*Creating a Hard Partition Block Flow*  
Fig. 38



*Creating a SubPartition Block Flow*

Fig. 39

	PMU Server	All MBMs/PBMs	The PMU gathers information about the system configuration.
	PMU Server	All MBMs/PBMs	The user associates a hard partition & sub partition with one or many EVs.
	PMU Server	All MBMs/PBMs	<b>MOVE EV7 TO PARTITION</b> stored in volatile database
	PMU Server	All MBMs/PBMs	The data is distributed to all MBMs and PBMs on the train with a FULL TRAIN transmission.
	PMU Server	All MBMs/PBMs	The PMU Server replies successfully after all the MBMs/PBMs have the update
	PMU Server	All MBMs/PBMs	The user associates each IO7 to a hard partition & sub partition
	PMU Server	All MBMs/PBMs	<b>ASSIGN IO7 TO SUB PARTITION</b> stored in volatile database
	PMU Server	All MBMs/PBMs	The data is distributed to all MBMs and PBMs on the train with a FULL TRAIN transmission.
	PMU Server	All MBMs/PBMs	The PMU Server replies successfully after all the MBMs/PBMs have the update.

*Creating a New Partition Flow Diagram (Part 1 of 2)*

Fig. 40

	PMU Server	All MBMs/PBMs	The user associates memory with the hard partition & sub partition
	PMU Server	All MBMs/PBMs	<b>ASSIGN MEMORY TO SUB PARTITION</b> stored in volatile database
	PMU Server	All MBMs/PBMs	The data is distributed to all MBMs and PBMs on the train with a FULL TRAIN transmission.
	PMU Server	All MBMs/PBMs	The PMU Server replies successfully after all the MBMs/PBMs have the update.
	PMU Server	All MBMs/PBMs	The user indicates that he/she is done setting up the partition.
	PMU Server	All MBMs/PBMs	<b>SAVE PARTITION ASSIGNMENTS</b>
	PMU Server	All MBMs/PBMs	The data is distributed to all MBMs and PBMs on the train with a FULL TRAIN transmission.
	PMU Server	All MBMs/PBMs	The PMU Server replies successfully after all the MBMs/PBMs have the update
	PMU Server	All MBMs/PBMs	Done. The user can now start the partition

*Creating a New Partition Flow Diagram (Part 2 of 2)*

Fig. 41

## RESET STATE

PMU	PMU Server	Group Leader	MBMs	CMMs	EV7s	After a group formation, the leader reconciles the partition database and starts up a partition coordinator for each partition. The PC is started on the lowest MBM id in the partition.
PMU	PMU Server	Part. Coord	MBM	CMMs	EV7s	The PC gathers information about the current state of the partition. Each MBM in the partition is queried for state of the CMMs with the <b>GET MBM CONFIGURATION</b> request
PMU	PMU Server	Part. Coord	MBMs	CMMs	EV7s	Each CMM in the partition is queried with the <b>GET CMM STATE</b> command.
PMU	PMU Server	Part. Coord	MBMs	CMMs	EV7s	Each CMM returns the status of the EV7s it controls.
PMU	PMU Server	Part. Coord	MBMs	CMMs	EV7s	The MBM assembles the data from all the CMMs into the reply to the <b>GET MBM CONFIGURATION</b>
PMU	PMU Server	Part. Coord	MBMs	CMMs	EV7s	Each PBM in the partition is queried for state of the CMMs with the <b>GET PBM CONFIGURATION</b> request
PMU	PMU Server	Part. Coord	PBMs	CMMs	EV7s	The PBMs reply with the IO7 configuration.
PMU	PMU Server	Part. Coord	MBMs	CMMs	EV7s	The P.C. sets the state of the partition to <b>RESET-IN-PROGRESS</b> with the <b>SET PARTITION STATE</b> .
PMU	PMU Server	Part. Coord	MBMs	CMMs	EV7s	The command goes to all MBMs in the entire system, not just the ones in this partition. The MBMs record the state. See the "MBM Start Partition State Diagram."
PMU	PMU Server	Part. Coord	MBMs	CMMs	EV7s	The Part. Coord. determines the PID of every EV7 in the hard partition and assigns a PID. The CMM is sent a <b>CONFIG RBOX/CBOX</b> packet with this PID value and defaults for router tables.
PMU	PMU Server	Part. Coord	MBMs	CMMs	EV7s	The CMM stores the PID information in ram that is accessible by the EV7 for later retrieval.
PMU	PMU Server	Part. Coord	MBMs	CMMs	EV7s	The Partition Coordinator performs a partition reset by issuing, individually, a <b>PULSED RESET</b> to each EV7 in the partition.
PMU	PMU Server	Part. Coord	MBMs	CMMs	EV7s	The CMM intercepts this command and performs the <b>PULSED RESET</b> on the EV7s. The CMM facilitates the FPGA load and SRAM load which then loads the XSROM the the <b>LOAD XSROM IMAGE</b> packet.

*Partition Start Flow Diagram (Reset State)*

Fig. 42

## DIAGNOSTIC STATE

PMU	PMU Server	Part. Coord	MBMs	CMMs	EV7s	The P.C. sets the state of the partition to <i>DIAGNOSTICS-IN-PROGRESS</i> with the <b>SET PARTITION STATE</b> .
PMU	PMU Server	Part. Coord	MBMs	CMMs	EV7s	The command goes to all MBMs in the entire system, not just the ones in this partition. The MBMs record the state. See the "MBM Start Partition State Diagram."
PMU	PMU Server	Part. Coord	MBMs	CMMs	EV7s	After all EV7's are in XSROM, the P.C. commands all EV7s individually to start the first diagnostic test. P.C. repeats this operation for each of the tests in the suite using the <b>EV7 START TEST</b> .
PMU	PMU Server	Part. Coord	MBMs	CMMs	EV7s	The CMMs receive the <b>EV7 START TEST</b> and pass the instructions up to the XSROM running on the EV7.
PMU	PMU Server	Part. Coord	MBMs	CMMs	EV7s	The suite of tests is started sequentially by the partition coordinator until all the tests are complete.
PMU	PMU Server	Part. Coord	MBMs	CMMs	EV7s	The Partition Coordinator asks the PMU server to perform the cable testing with a <b>RECONFIGURE CABLING</b> .
PMU	PMU Server	Part. Coord	MBMs PBMs	CMMs	EV7s	The PMU Server coordinates the cable testing and issues <b>GET MBM IP CABLING</b> and <b>GET PBM IO CABLING</b> . The targets then issue <b>RECEIVE CABLE ID</b> and <b>SEND CABLE ID</b> .
PMU	PMU Server	Part. Coord	MBMs	CMMs	EV7s	When the cable test operation is complete, the PMU Server responds to the requesting P.C. with a completion status. The P.C. gets the latest results from the PMU Server with the <b>GET CABLE CONFIGURATION</b> request.
PMU	PMU Server	Part. Coord	MBMs	CMMs	EV7s	The PMU Server supplies the latest cable configuration data in the reply.
PMU	PMU Server	Part. Coord	MBMs	CMMs	EV7s	The cable connectivity, the diagnostic test results, the striping information, and the partition database are used in CTABLE to calculate PIDs and Routing Tables.

Figure 43 Partition Start Flow Diagram, Diagnostic State

## CONFIGURE ROUTER

	PMU Server	Part. Coord	MBMs	CMMs	EV7s	The P.C. sets the state of the partition to <b>CONFIGURE-ROUTER-IN-PROGRESS</b> with the <b>SET PARTITION STATE</b> .
	PMU Server	Part. Coord	MBMs	CMMs	EV7s	The command goes to all MBMs in the entire system, not just the ones in this partition. The MBMs record the state. See the "MBM Start Partition State Diagram."
	PMU Server	Part. Coord	MBMs	CMMs	EV7s	<b>QUIESCE PARTITION</b> is sent to all EV7s in the partition. In this example, the EV7s are in XSROM, which is already a quiesced state.
	PMU Server	Part. Coord	MBMs	CMMs	EV7s	This command is serviced by the CMM and performs the necessary steps to have the EV7 (still running XSROM) configure the router.
	PMU Server	Part. Coord	MBMs	CMMs	EV7s	<b>CONFIG RBOX/CBOX</b> is sent to all EV7s in the partition
	PMU Server	Part. Coord	MBMs	CMMs	EV7s	This command is serviced by the CMM and performs the necessary steps to have the EV7 (still running XSROM) configure the router.
	PMU Server	Part. Coord	MBMs	CMMs	EV7s	Continued on next flow

*Partition Start Flow Diagram (Configure Router)*

Fig. 44

**CONTINUE PARTITION  
and  
PARTITION RUNNING**

	PMU Server	Part. Coord	MBMs	CMMs	EV7s	The P.C. sets the state of the partition to <b>CONTINUE-PARTITION-IN-PROGRESS</b> with the <b>SET PARTITION STATE</b> .
	PMU Server	Part. Coord	MBMs	CMMs	EV7s	The command goes to all MBMs in the entire system, not just the ones in this partition. The MBMs record the state. See the "MBM Start Partition State Diagram."
	PMU Server	Part. Coord	MBMs	CMMs	EV7s	The non-primary EV7s are given an <b>EV7 START TEST</b> that sets them waiting for a flag to jump into the running image.
	PMU Server	Part. Coord	MBMs	CMMs	EV7s	This command is forwarded by the CMM and EV7s are now ready and waiting
	PMU Server	Part. Coord	MBMs	CMMs	EV7s	The partition is now completely configured. All EV7s are running XSROM. The P.C. now directs the primary EV7 to start the SRM console via the <b>LOAD IMAGE</b>
	PMU Server	Part. Coord	MBMs	CMMs	EV7s	The CMM intercepts this command and performs the necessary operations to get the Primary EV7 loaded and running.
	PMU Server	Part. Coord	MBMs	CMMs	EV7s	The Primary EV7, writes a flag in each of the other EV7s , indicating that they should jump into the running image.
	PMU Server	Part. Coord	MBMs	CMMs	EV7s	The P.C. sets the state of the partition to <b>PARTITION RUNNING</b> with the <b>SET PARTITION STATE</b> .
	PMU Server	Part. Coord	MBMs	CMMs	EV7s	The command goes to all MBMs in the entire system, not just the ones in this partition. The MBMs record the state. See the "MBM Start Partition State Diagram."
	PMU Server	Part. Coord	MBMs	CMMs	EV7s	The P.C. notifies the PMU server in the system that the partition has been started.
	PMU Server	Part. Coord	MBMs	CMMs	EV7s	The PMU Server notifies all the connected PMUs that the partition has been started.
	PMU Server	Part. Coord	MBMs	CMMs	EV7s	The End

*Partition Start Flow Diagram (Running)*

Fig. 45

OS via SRM	PMU Server	Part. Coord HP#2	MBMs PBMs	CMMs in HP#2	EV7s	OS uses SRM callback to issue <b>ADD EV7 TO RUNNING PARTITION</b> listing EV7#4 and EV7#5 and Hard Partition #2 Sub #1. The PMU Server handles the request.
OS via SRM	PMU Server	Part. Coord HP#2	MBMs PBMs	CMMs in HP#2	EV7s	The PMU Server knows the address of the partition coordinator and forwards the <b>ADD EV7 TO RUNNING PARTITION</b> to the appropriate part. coord.
OS via SRM	PMU Server	Part. Coord HP#2	MBMs PBMs	CMMs in HP#2	EV7s	The PMU Server issues a <b>SET PARTITION STATE</b> to <b>PARTITION CHANGE IN-PROGRESS</b> to all MBMs via the train mechanism.
OS via SRM	PMU Server	Part. Coord HP#2	MBMs PBMs	CMMs in HP#2	EV7s in HP#2	All MBMs/PBMs track this state in their replicated data base. PMU Server access is limited until the partition change completes.
OS via SRM	PMU Server	Part. Coord HP#2	MBMs PBMs	CMMs in HP#2	EV7s in HP#2	The Part. Coord. directs the PMU Server to <b>RECONFIGURE CABLING</b> to get the current status of the cables
OS via SRM	PMU Server	Part. Coord HP#2	MBMs PBMs	CMMs in HP#2	EV7s in HP#2	The PMU Server coordinates the cable testing. Details are in a later section of this document. Results are stored at the PMU Server. <b>GET MBM IP CABLING</b> , <b>GET PBM IO CABLING</b> , <b>SEND CABLE ID</b> , <b>RECEIVE CABLE ID</b> .
OS via SRM	PMU Server	Part. Coord HP#2	MBMs PBMs	CMMs in HP#2	EV7s in HP#2	Upon successful completion of the cable test, the Part. Coord. issues a <b>GET CABLING CONFIGURATION</b> to the PMU Server to get the latest cable data.
OS via SRM	PMU Server	Part. Coord HP#2	MBMs PBMs	CMMs in HP#2	EV7s in HP#2	The PMU Server responds with the current cable connectivity data base.
OS via SRM	PMU Server	Part. Coord HP#2	MBMs PBMs	CMMs in HP#2	EV7s in HP#2	The Part. Coord validates the cable connectivity. If okay <b>ADD EV7 TO RUNNING PARTITION</b> is distributed to all MBMs/PBMs on the LAN via the train mechanism.
OS via SRM	PMU Server	Part. Coord HP#2	MBMs PBMs	CMMs in HP#2	EV7s in HP#2	All MBMs/PBMs modify their replicated data bases to track this change to the free pool and to partition #2.
OS via SRM	PMU Server	Part. Coord HP#2	MBMs PBMs	CMMs in HP#2	EV7#4 EV7#5	<b>PULSE EV7 RESET</b> is directed at EV7#4 and EV7#5 (command issued to the CMM)
OS via SRM	PMU Server	Part. Coord HP#2	MBMs PBMs	CMMs in HP#2	EV7#4 EV7#5	EV7#4 and EV7#5 are reset. The CMM facilitates the SROM load. The XSROM is started with the <b>LOAD IMAGE</b> .
OS via SRM	PMU Server	Part. Coord HP#2	MBMs PBMs	CMMs in HP#2	EV7s in HP#2	continues on next figure


Add EV7 Flow Diagram (Part 1)

Fig. 46 A

OS via SRM	PMU Server	Part. Coord HP#2	MBMs PBMs	CMMs in HP#2	EV7s in HP#2	The Part. Coord. calculates the routability of the proposed new configuration. If there is a routing error the command finishes with error. When routing is successful, a SET PARTITION STATE to CONFIG-ROUTE-IN-PROGRESS
OS via SRM	PMU Server	Part. Coord HP#2	MBMs PBMs	CMMs in HP#2	EV7s in HP#2	All MBMs receive the SET PARTITION STATE via the train mechanism and set their databases to indicate that HP#2 is in the CONFIG-ROUTE-IN-PROGRESS state
OS via SRM	PMU Server	Part. Coord HP#2	MBMs PBMs	CMMs in HP#2	EV7s in HP#2	When the train completes, the Part. Coord. commands each EV7 in hard partition #2, individually, to QUIESCE.
OS via SRM	PMU Server	Part. Coord HP#2	MBMs PBMs	CMMs in HP#2	EV7s in HP#2	QUIESCE is received at the CMMs in the partition and they take actions to quiesce the EV7. EV#4 and EV#5 are in XSROM and do not have to be quiesced.
OS via SRM	PMU Server	Part. Coord HP#2	MBMs PBMs	CMMs in HP#2	EV7s in HP#2	CONFIG RBOX/CBOX is sent to all EV7s in the partition, which includes EV#4 and EV#5 which are in XSROM.
OS via SRM	PMU Server	Part. Coord HP#2	MBMs PBMs	CMMs in HP#2	EV7s in HP#2	CONFIG RBOX/CBOX is received at the CMMs and they direct the EV7 to config the RBOX and the CBOX.
OS via SRM	PMU Server	Part. Coord HP#2	MBMs PBMs	CCMs in HP#2M	EV7s in HP#2	SET PARTITION STATE to CONTINUE-PARTITION-IN-PROGRESS. This is distributed via the train mechanism.
OS via SRM	PMU Server	Part. Coord HP#2	MBMs PBMs	CCMs in HP#2M	EV7s in HP#2	All MBMs/PBMs track the state of the partition in their replicated data base.
OS via SRM	PMU Server	Part. Coord HP#2	MBMs PBMs	CCMs in HP#2M	EV7s in HP#2	CONTINUE PARTITION is sent to all EV7s in the partition, including EV7#4 and EV7#5.
OS via SRM	PMU Server	Part. Coord HP#2	MBMs PBMs	CCMs in HP#2M	EV7s in HP#2	All MBMs/PBMs track the state of the partition in their replicated data base.
OS via SRM	PMU Server	Part. Coord HP#2	MBMs PBMs	CCMs in HP#2M	EV7s in HP#2	SET PARTITION STATE to PARTITION RUNNING is issued on the train.
OS via SRM	PMU Server	Part. Coord HP#2	MBMs PBMs	CCMs in HP#2M	EV7s in HP#2	All MBMs/PBMs track the state of the partition in their replicated data base.
OS via SRM	PMU Server	Part. Coord HP#2	MBMs PBMs	CCMs in HP#2M	EV7s in HP#2	The Part. Coord. notifies the PMU Server that the ADD EV7 TO PARTITION is done.
OS via SRM	PMU Server	Part. Coord HP#2	MBMs PBMs	CCMs in HP#2M	EV7s in HP#2	The PMU Server responds to the original ADD EV7 to PARTITION.
OS via SRM	PMU Server	Part. Coord HP#2	MBMs PBMs	CCMs in HP#2M	EV7s in HP#2	The End.


Add EV7 Flow Diagram (Part 2)

Fig. 46B




*Add vs Move*

Fig. 47



*Destroying a soft partition*

Fig. 48





| *Destroying a hard partition*

Fig. 49


This flow is an example of SM protocol activity when a CPU module fails. The entire configuration is one partition, and the failing CPU module is not the primary.

The operating system does crash, and is restarted. The failed CPU module is replaced and the original complete configuration is restored.



*EV7 Failure/Replace Flow Diagram (Part 1)*

Fig. 50




*EV7 Failure/Replacement Flow Diagram (Part 2)*

Fig. 51

 PMU	MBM1	MBM2	PBM0	New MBM3	The Marvel system is running with MBM1, MBM2, and MBM3. A new 8P is rolled up and plugged into the SM LAN and powered up.
 PMU	MBM1	MBM2	PBM0	New MBM3	MBM3 recognizes that it is joining a already formed group of processors and has a different membership list than the group it has joined. MBM3 goes into a passive listening state.
 PMU	MBM1	MBM2	PBM0	New MBM3	The New MBM3 is not defined as a member in the old group, so it is isolated from the group.
 PMU	MBM1	MBM2	PBM0	New MBM3	The operator issues a <b>SET MEMBERSHIP CONFIGURATION</b> to the PMU to change the membership list from {MBM1,MBM2,PBM0} to {MBM1,MBM2,MBM3,PBM0}. The PMU broadcasts this to the entire LAN.
 PMU	MBM1	MBM2	PBM0	New MBM3	All LAN members receive this membership list and change their expected membership data. Then they all participate in new group formations.
 PMU	MBM1	MBM2	PBM0	New MBM3	The group is formed MBM3 is now an active member.

*Set Membership Configuration Flow Diagram*

Fig. 52



*IP Cable Configuration Block Diagram*

**Fig. 53**

The EV7 IDs (x,y) are determined by the thumb-wheel setting using the following algorithm:

$$x \text{ (E,W coordinate)} = ((\text{Rack Number} \gg 2) * 8) + ((\text{MBM number} \gg 1) * 4) + \text{CMM number}$$

$$y \text{ (N,S coordinate)} = ((\text{Rack Number} \& 0x03) * 4) + ((\text{MBM number} \& 0x01) * 2) + \text{EV7 number}$$

where:

Rack Number is the high order half byte of the MBM thumb-wheel

MBM Number is the low order half byte of the MBM thumb-wheel


CMM Number is from 0 to 3 within an MBM

EV7 number is 0 or 1 within a CMM

In a similar manner when the x,y axis coordinates of an EV7 are known, the thumb-wheel numbers can be derived and inserted into the IP address for the MBM, CMM and EV7s.


*EV7 Coordinate addressing relationship to thumbwheel addressing*

Fig. 54




*IO7 Cabling Block Diagram*

Fig. 55




*Get Cable Configuration Block Diagram*

Fig. 56




*IP Cable Addition/Deletion Block Diagram*

Fig. 57




*PMU Cabling Assistant Block Diagram*

**Fig. 58**



*Proxy Forwarding Block Diagram*

Fig. 59



*Virtual Console Terminal Overview*

Fig. 60

CMM	EV7	COM1 Port	COM2 Port
1	1	323	324
1	2	325	326
2	1	327	328
2	2	329	330
3	1	331	332
3	2	333	334
4	1	335	336
4	2	337	338

*Virtual Terminal Telnet*

Fig. 61


## Virtual Terminal Session

	PMU Server	Telnet Srv for Pri EV7	MBMs	CMMs	EV7s	The PMU platform is plugged into an available port on the SM LAN
	PMU Server	Telnet Srv for Pri EV7	MBMs	CMMs	EV7s	The PMU sends a Get Telnet IP Address/Port request to PMU Server for the address of the telnet server for the primary EV7 for a partition
	PMU Server	Telnet Srv for Pri EV7	MBMs	CMMs	EV7s	PMU Server gives to the PMU the address of the telnet server for the primary EV7 <b>(Get Telnet IP Address/Port response packet)</b>
	PMU Server	Telnet Srv for Pri EV7	MBMs	CMMs	EV7s	User starts a telnet session on the PMU using the telnet server address from previous step
	PMU Server	Telnet Srv for Pri EV7	MBMs	CMMs	EV7s	Telnet server passes characters to Virtual Console task on MBM for Primary EV7
	PMU Server	Telnet Srv for Pri EV7	MBMs	CMMs	EV7s	MBM for Primary EV7 passes characters to CMM for Primary EV7 using <b>PUT_CHAR</b> packet
	PMU Server	Telnet Srv for Pri EV7	MBMs	CMMs	EV7s	CMM strips off IP addressing info and passes the character to Primary EV7 through the virtual console uart registers
	PMU Server	Telnet Srv for Pri EV7	MBMs	CMMs	EV7s	Primary EV7 responds with characters back to CMM
	PMU Server	Telnet Srv for Pri EV7	MBMs	CMMs	EV7s	CMM envelops character in <b>PUT_CHAR</b> protocol packet and gives to MBM
	PMU Server	Telnet Srv for Pri EV7	MBMs	CMMs	EV7s	MBM Virtual Console task passes character to Telnet server
	PMU Server	Telnet Srv for Pri EV7	MBMs	CMMs	EV7s	Telnet server passes character to Telnet session on PMU

*Virtual Terminal Flow Diagram*


Fig. 62

## SET BASE TIME



*Set Base Time Flow Diagram*

Fig. 63



*DHCP Block Diagram*

Fig. 64

<i>Content</i>	<i>TFTP Filename</i>	<i>Block Size</i>	<i>Sector Start(offset)</i>	<i>Sector End</i>
MBM/PBM Firmware	"mbmfw", "pbmfw"	2 MB	0	7
CMM, CMM_FSL, FPGA, SROM, XSROM Firmware	"cmmfw", "cmmfsl", "cmmpgfa", "sromfw", "xsromfw"	0.5 MB	8(0,tbd,tbd,tbd,tbd,tb d)	9
Error Logs		1 MB	10	13
MBM/PBM FSL Firmware	"mbmfsl", "pbmfsl"	0.5 MB	14	15
NVRAM – partition database		0.75 MB	16	18
SRM Firmware	"srmfw"	2 MB	20	27
FPGA loaded by PBM on PCI Drawers	"pbmfpga"	0.25MB	28	28
MBM/PBM Boot[if required by HW]		0.25	19(0,0x30000)	19

*Flash Layout*

*Fig. 65*

Byte Offset	Value
0.. 7	0x010100005500aaff
8.. 19	Image Revision in ASCII
20.. 23	Vendor String in ASCII (CPQ)
24.. 31	Module ID in ASCII  (SRMFW,MBMFW,MBMFSL,CMMFW,CMMFSL,SROMFW,XSROM FW,CMMFPGA)
32.. 35	Firmware Type in ASCII (ALPH, X86)
36.. 43	0x00
44.. 47	Code Length in bytes
48.. 59	ROM Object Name (FW,FSL,SROM,XSROM,FPGA)
60.. 63	0x11223344


*Image Header*

Fig. 66

All MBMs PBMs	MBM	CMM	SRM	The primary CPU in the partition issues the command >>> set bootdef_dev dka100 This becomes a <b>STORE ENVIRONMENT VARIABLES</b> packet
All MBMs PBMs	MBM	CMM	SRM	MBM waits for the train to arrive and then retransmits the <b>FULL TRAIN MESSAGE</b> with the <b>STORE ENVIRONMENT VARIABLES</b> payload.
All MBMs PBMs	MBM	CMM	SRM	All MBMs receive the train, copy out the payload and hold that payload as a pending command. The MBM passes along the <b>FULL TRAIN MESSAGE</b> , and wait for the empty train.
All MBMs PBMs	MBM	CMM	SRM	The train makes it full circle back to the originating MBM. It then commits the SRM environment variables to their flash. It then puts the <b>FULL TRAIN MESSAGE</b> into the train payload and reissues it.
All MBMs PBMs	MBM	CMM	SRM	All MBMs commit the data to their flashes and pass on the <b>FULL TRAIN MESSAGE</b> .
All MBMs PBMs	MBM	CMM	SRM	The train makes it full circle back to the originating MBM. It now can respond to the <b>STORE ENVIRONMENT VARIABLES</b> command and sends out the <b>EMPTY TRAIN MESSAGE</b> .
All MBMs PBMs	MBM	CMM	SRM	The reply passes through the CMM and back to the EV7 running the SRM console with a successful completion status.


*SRM Environment Vars Flow Diagram*

Fig. 67



*Firmware Load and Upgrade Block Diagram*

Fig. 68



*Upgrading CMM Firmware Flow Diagram*

Fig. 69

ERROR LOG ENTRY											
Size (dec)	Start (hex)	End (hex)	Bit 7	Bit 6	Bit 5	Bit 4	Bit 3	Bit 2	Bit 1	Bit 0	
2	0	1	Entry Number (1-0xffff before wrapping)								
6	2	7	Time Stamp (ssmmhhDDMMYY)								
2	8	9	Entry Size								
n	A	n+A	Entry Data								

*Error Log Entry Format*

Fig. 70

ERROR ENTRY DATA																		
Size (dec)	Start (hex)	End (hex)	Bit 7	Bit 6	Bit 5	Bit 4	Bit 3	Bit 2	Bit 1	Bit 0								
4	0	3	Entry Number															
1	4	4	Severity Level															
1	5	5	Entity in Error															
1	6	6	Instance															
2	7	8	Error Code															
16	9	18	Serial Number															
n	19	n+19	Variable Data															
<b>Entry Number</b> IP address of CMM, MBM, PBM, or EV7 that encountered the error.																		
<b>Severity Level</b> Informational=0; Warning=1; Error=2;																		
<b>Entity in Error</b> The device code for the device in error (e.g. CMM,RIMM, EV7, Thermal, Volatage)																		
<b>Instance</b> The instance of the entity																		
<b>Error Code</b> Error enumeration or index into a set of text messages																		
<b>Serial Number</b> Identifying address where the error occurred.																		
<b>Variable Data</b> Additional data specific to this error code.																		


*Error Entry Data Format*

Fig. 71

	PMU Server	MBM	CMM	EV7	ERROR REPORTING to MBM with PMU Alert, MBM OCP
	PMU Server	MBM	CMM	EV7	MBM places error entry to OCP, logs error to NVRAM, responds positively to originating CMM and sends ERROR REPORTING to primary EV7 of affected partition, and also to the PMU Server.
	PMU Server	MBM	CMM	EV7	The CMM acknowledges receipt of the ERROR REPORT to the MBM CMM sends SYSEVENT to interrupt OS and sends the error report message. The PMU Server sends an ERROR ALERT to all the PMUs.
	PMU Server	MBM	CMM	EV7	The OS sees an interrupt, retrieves the SYSEVENT, and acknowledges the interrupt. The MBM receives the ERROR REPORT response from the CMM.
	PMU Server	MBM	CMM	EV7	The PMU receives the ERROR ALERT. The CMM receives the response from the SYSEVENT

*Error Reporting Flow Diagram*

Fig. 72



Field Replaceable Unit

Fig. 73

<i>Line</i>	12345678901234567890
1	<Overall Progress>
2	<Current State>
3	<Location within State>
4	<Error Message>

*OCP Template*

Fig. 74

0	1	3	4	5	6	X
P		PP	F	P		
1	1	0	1	1	1	1
EV 5 RIMM 2 Parity						


*OCP 8P Example*

Fig. 75

0	1	3	4	5	6	X
P		PP	F	P		
1	1	0	1	1	1	1
Power      Halt						
Reset						

*OCP Button Label Example*

Fig. 76



*OCP Switches Block Diagram*

Fig. 77

Command	Purpose	Process
SET ATTENTION INDICATOR	Light, Extinguish LED at MBM, PBM, CMM or Cabinet	Take action according to the desired state in the request.
SET KNOB	Name/value pair to control some MBM/PBM capabilities	Coded into the Image is a list of variable names that can be modified.
GET KNOB	The value currently maintained for the named variable is returned.	Check to see if the name is in the list and return it's current value in the response.
READ	Allow debug read of physical memory or I/O.	Check on validity of request and if valid, read physical memory space or direct I/O space. For debug.
WRITE	Allow debug write of physical memory or I/O.	Check on validity of request and check MMU protection privileges to write in the space so as not to cause protection violations. If not protected, write the block. For debug.

*Miscellaneous Command Handling*

Fig. 78

Command	Arg Count	Arguments	Result	Error Handling
Show_config	0		See Section "Show Configuration with FRU Data"	Make requests to PMU Server
Reset	2	1-Partition No 2-Sub Partition No	Returns OK or error	Send Reset Partition to PMU Server
Power_on	2	1-Partition No 2-Sub Partition No	Returns OK or error	Send Power On Partition to PMU Server
Power_off	2	1-Partition No 2-Sub Partition No	Returns OK or error	Send Power Off Partition to PMU Server
Halt_on	2	1-Partition No 2-Sub Partition No	Returns OK or error	Send Halt on Partition to PMU Server
Halt_off	2	1-Partition No 2-Sub Partition No	Returns OK or error	Send Halt Off Partition to PMU Server
Prepare_EV7_List	2-16	1-MBM Rack-thumb-wheel, 2-Ev7 Id (0-7) up to 8 pairs	OK if all elements are in the same hard partition or free pool	Saves this value in MBM RAM for use with the next Add EV7s, Free EV7s. Lasts until next Prepare EV7 List.
Add_Ev7s	2	1-Partition No 2-Sub Partition No	Take the values in the Ev7 list and add it to the partition.	Send command to PMU Server.
Free_Ev7s	0		Take the values in the Ev7 list and remove them from the partition indicated.	Send command to PMU Server.
Save_partition	2	1-Partition No 2-Sub Partition No	The partition database gets stored to NVRAM.	Send command to PMU Server.
Destroy_partition	2	1-Partition No 2-Sub Partition No	Reset & Free all Ev7s from partition.	Send command to PMU Server.
Ev7_test	3	1-MBM Rack-thumb-wheel, 2-Ev7 Id (0-7), 3-Test Number	Test Status	Send a Ev7 Start Test request to the CMM.
Add_cable	5-6	1-Source - MBM Rack-thumb-wheel, 2-Ev7 Id (0-7), 3-Port(N,S,E,W), 4-Destination - MBM/PBM Rack-thumb-wheel, 5- Ev7 Id(0-7) or IO7 Id(0-3), 6-Port(N,S,E,W) or blank when IO7	This command assists in locating the proper connector pair to connect the cable. The leds at each connector are lit until the connection is complete.	The commands Set Cable Test Signal State and Get Cable Test Signal State are sent to the appropriate MBM and PBM to cause the leds to light and check the connection itself.
New_cabling	0		Redo cabling tests.	Send Reconfigure Cabling to PMU Server.
Show_cabling	0		Displays a list of IP & IO Cabling	Send Retrieve Cabling Configuration to PMU Server.
Virt_console	3	1-Partition No, 2-Sub Partition No, 3-COM Port(1,2)	Open a session with primary EV7 & intercept COMx Port Data.	Use PutChar streams for both display and keyboard data until the keyboard data sequence 'ESC'ESC'S'M' is recognized as an exit of the session.
Get_fans	1	1-MBM/PBM Rack-thumb-wheel	RPM and threshold that fans are running at.	Determine appropriate IP address for destination and send a Get Fan RPM Speed message.
Set_fan	3	1- MBM/PBM Rack-thumb-wheel, 2-Instance of Fan, 3-RPM	Error or OK	Determine appropriate IP address for destination and send a Set Fan RPM Speed message.
Error_counts	0		Returns a list of the error counts on all MBM/PBM error logs.	Send Error Log Count request to each MBM/PBM.
Error_clear	1	1- MBM/PBM Rack-thumb-wheel	OK	Send Error Log Clear request to destination
Get_errors	1	1-count of the number of errors to be reported on each device.l	A list of the last Error Messages in English as it would appear on the OCP with any qualifying data formatted as appropriate. This is repeated for each MBM and PBM.	Get the ERROR LOG COUNT from each micro. Send Error Log Entry Retrieval Requests to each micro using the highest number as the 1 <sup>st</sup> request.

Fig. 79A

Command	Arg Count	Arguments	Result	Handling
Get_time	0		Date and time is displayed as dd/mm/yy hh:mm:ss	Use Get base Time command
Set_time	1	1-Date and Time entry in format: "dd/mm/yy-hh:mm:ss"	Redisplays date and time	Set Base Time and Announce Base Time Change is sent to all MBMs.
Req_knobs	2-3	1- MBM/PBM Rack-thumb-wheel, 2- Device("MBM", "PBM", "CMM"), 3- If CMM in 2, then number (0-3)	Names and values of all possible knobs are listed.	Request Knob command for all possible knobs for that device.
Set_knob	4-5	1- MBM/PBM Rack-thumb-wheel, 2- Device("MBM", "PBM", "CMM"), 3-If CMM in 2, then number (0-3) 2-Knob name, 3-Knob Value	OK	Set Knob on requested micro.
Firmware_version	3-4	1- MBM/PBM Rack-thumb-wheel, 2- Device("MBM", "PBM", "CMM"), 3-If CMM in 2, then number (0-3) 4-Module ("CMM", "FPGA", "SRAM", "XSROM", "MBM", "PBM", "CMM_FS_L", "MBM_FSL", "PBM_FSL", "PBM_FPGA")	Returns the version number	Report Firmware Version Command
Firmware_upgrade	4-5	1- MBM/PBM Rack-thumb-wheel, 2- Device("MBM", "PBM", "CMM"), 3-If CMM in 2, then number (0-3) 4-Module ("CMM", "FPGA", "SRAM", "XSROM", "MBM", "PBM", "CMM_FS_L", "MBM_FSL", "PBM_FSL", "PBM_FPGA") 5-TFTP Server IP Address	Makes the terminal into a PPP serial link making TFTP requests until completion of the transfer or a timeout occurs. Possible return values are: "Complete", "Timeout", "File too long", "File too short"	Send the Upgrade Firmware request to the MBM or PBM or CMM. Make the backup copy where required. MBM has a copy of CMM program.

Fig. 79 B

Command	Arg Count	Arguments	Result	Handling
Load_test_version	4-5	1- MBM/PBM Rack-thumb-wheel, 2- Device("MBM", "PBM", "CMM"), 3-If CMM in 2, then number (0-3) 4-Module ("CMM", "FPGA", "SRAM", "XSROM", "MBM" " "PBM", "CMM_FSL", "MBM_FSL", "PBM_FSL", "PBM_FPGA") 5-TFTP Server IP Address	Makes the terminal into a PPP serial link making TFTP requests until completion of the transfer or a timeout occurs. Possible return values are: "Complete", "Timeout", "File too long", "File too short"	Send the Load Test Version Command to the MBM or PBM or CMM, which maintains a copy of the program in MBM memory.
Disable_test_version	3-4	1- MBM/PBM Rack-thumb-wheel, 2- Device("MBM", "PBM", "CMM"), 3-If CMM in 2, then number (0-3) 4-Module ("CMM", "FPGA", "SRAM", "XSROM", "MBM" " "PBM", "CMM_FSL", "MBM_FSL", "PBM_FSL", "PBM_FPGA")	OK	Send Disable_test_version command

*CLI Commands*

Fig. 79 C

Knob Name	Possible Values	Usage
CLI_PORT_SPEED	1200, 2400, 4800, 9600, 19200, 38400, 57600, 115200, 230400	Speed between COM Port and Modem (default 57600bps)
CLI_DATA_BITS	8, 7	COM Port UART uses 7 or 8 data bits before stop (default 8)
CLI_STOP_BITS	1, 1.5, 2	COM Port UART sends 1, 1.5 or 2 stop bits to modem (default 1)
CLI_FLOW_CTL	HW, SW, NONE	Flow Control: HW – RTS/CTS signals (default), SW-Xon/Xoff bytes, None
CLI_MODEM	YES, NO	If no modem is connected use no; otherwise the following set of CLI_MDM knobs are required.
CLI_MDM_INIT	AT string for modem initialization	On each hang-up or drop of carrier signal this command is sent to the modem (default is "ATZ"). If modem doesn't respond with OK, 3 retries are attempted.
CLI_MDM_DIAL	AT string when we dial out to drop an alert message or dial-back.	Prefix for dialing the number indicated in the alert number or dial-back number. (default is "ATDT"). If modem responds with OK, communication is considered to be established, else 3 retries are attempted.
CLI_DIAL_BACK	Phone Number with dialing pauses	For security purposes whenever a connection is made, the program will hangup and dial the indicated number to establish connection. (default empty)
CLI_DIAL_ALERT	Phone Number with possible dialing pauses for a receiver of alert messages.	If an error message has an alert indication, the text portion that would be formatted for the OCP is sent after establishing a modem connection with the indicated number. Note: There is no paging support TAPI, Alphanumeric or Numeric implied by this option.
CLI_PASSWORD	The only password that is allowed to at time of connection. A password prompt is used.	The password prompt appears on a modem connection requesting the password entry to continue. The entry must match the null terminated string belonging to this knob. If not, a hang-up command is sent to the modem. A default password will exist <TBD>, if none has been assigned.

#### *Modem Knobs*

Fig. 80

Capability	Ramification during MBM failure
Proxy Forwarding	LAN messages to the CMM fail.
CMM presence watchdog	There is no notification of failure
TFTP Server	CMM is unable to be firmware upgraded
Error repository & distribution	CMM is unable to log errors
Power hierarchy	Partition power changes cannot occur
Time Services	CMM does not receive regular time updates
Virtual Console Terminal	Console input/output unavailable if primary in path
SRM Env. Var. Repository	Env. Writes fail; SRM write callbacks fail.

MODEM KNOBS FOR CONNECTION  
TO CLI PORT

Fig. 81

Command	OS / SRM Up	All EV7s	Maj. Grp.	Comment
Distribute Partition Database with EV7 change, Add EV7, Delete EV7	T	T	T	Provided the new EV7 is still in the majority group and routable.
Set Partition Delta Time, Store Environment Variables, Assign Sub partitions, Assign Memory & IO7	X	X	T	This affects the database; but when joined back again majority wins.
Reset, Halt, Quiesce, Continue	T	T	X	This doesn't affect the non-volatile database.
Power On, Power Off, Change Primary	X	X	X	Never allowed on a split partition

*Operation Limitations in a degraded system*

Fig. 82

Operations	Structures	Distribution Commands	RAM	NVRAM
The LAN Group protocol messages use either broadcast or request/response messages in forming a group.	GroupId(RAM), MajGID(RAM), MicroprocessorSet(NVRAM)	NewGroup, Accept, Join, SetMembership	x	x
The Group Leader, after forming a new group, checks the copy of all members' partition databases and distributes the copy. EV7 changes noted by the PMU Server are distributed via a partition database change. The partition coordinator, when reconfiguring the partition, may need to distribute commands that change the volatile copy of the partition database.	EV7/IO7, memory locations, assignments and status.	Request partition database, Distribute partition database	x	x
The PMU Server initialization includes the cable testing and distribution of that volatile database information	Cable Configuration	Retrieve Cabling Configuration	X(not distributed)	
The DHCP Server distributes the list of DHCP leases and their changes.	DHCP Leases	Distribute DHCP Lease Data	x	x
The Partition Coordinator distributes to all members any changes in the status of the partition's state and attributes to allow fail-over recovery to another partition coordinator. Changes to Memory, IO7 and Community assignments among the sub-partitions. Partition states are affected by the ongoing starting, routing, loading, halting, resetting and power controls on the entire partition. Attribute fields on a partition are: 1) OS Watchdog Interval and Action Mask, 2) BB_Watch Delta Time, 3) SRM Environment Variables.	Partition State and attribute fields(RAM). The state and OS Watchdog is kept in volatile RAM and all other attributes are stored in NVRAM location for that partition.	Set Partition State and attributes or Distribute One Partition's Database	x	x
MBM maintains in RAM status on its own CMMs, EV7s and memory. I2C data, error log count, Knobs and OCP data is maintained in RAM.	CMM EV7 Status, I2C sensor and EEPROM values, volatile Knobs.	MBM Report Configuration	X(not distributed)	
MBM maintains the current use of the OCP switch settings and some Knobs. Default is entire system enclosure.	OCP switch control, permanent Knobs.	OCP to Partition Assignment	X(not distributed)	X(not distributed)
PBM maintains a RAM copy of its own IO7 Ids, I2C data and error log count.	IO7 riser Ids, I2C sensor and EEPROM values.	PBM Report Configuration, Receive Cable ID	X(not distributed)	
PBM receives from SRM the PCI Configuration Data and retrieves it on request. MBM, PBM and CMM knob data are kept in NVRAM to tailor behavior.	PCI Config Space(RAM), Knob Data.	Store PCI Slot Info, Read PCI Slot Info, Set Knob Data, Get Knob Data	X(not distributed)	X(not distributed)

*Data Base Grouping*

Fig. 83

84.000

84,002 IN FLASH

NON-VOLATILE PARTITION DB

N/S	E/W	HARD	SUB
0	0	7	
.	.	.	
.	.	.	
.	.	.	
x	y	z	

x:0-15 y:0-15 z:0-255

VARIABLE SIZE; DENSE;

(R)

MICROPROCESSOR SET

RACK	BOX
1	P 1
2	17
...	...

~ 84,004

P - PROCESSOR

I - I/O

NON-VOLATILE

RANGE: RACK: 0-32  
BOX: 0-32

VARIABLE LENGTH  
MAX SIZE:

(R) REPLICATED  
SYSTEM  
WIDE

SRM ENV VARs

HARD 0 SUB 0

bootdef\_dev

ewa0\_mode

.....

~ 84,006

PER SUB PARTITION

NON-VOLATILE

FIXED 2KB

(R)

~ 84,008

DELTA TIME

(R)


PER SUB PARTITION

NON-VOLATILE

FIXED: 6 BYTES x 256

REPLICATED DATABASE

FIG. 84A







Fig. 85




PBM Hardware Overview

Fig. 86


Errors Codes	Device Type	Reporters	OCP/CLI Text	OS	Alert	Sysevent	EEROM	Loc Data	System	OCP
1	none	Anyone	Free Text	x	?	?		x	?	
2	none	Anyone	Binary with no text	x	?	?		x	?	
1	EV7	CMM	EV7 xx Failure Code xx on running OS	x	x	x		x	x	
2	EV7	CMM	EV7 xx Failure in Routing to EV7 xx	x	X					
3	EV7	CMM	EV7 xx Power On Failed	x	x				x	
4	EV7	CMM	EV7 xx Failed on Test xx with Status xx	x	x					
5	EV7	CMM	EV7 xx Overheated at xxx Fahr.	x	x	x			x	
6	EV7	CMM	EV7 xx xx.x Power at xx.x	x	x	x			x	
7	RIMM	CMM	Memory Test xx Failure for RIMM xx	x	x			x	x	
8	FPGA	CMM	CMM x FPGA Failure xx	x	x				x	
9	CMM	CMM	CMM x POST Error xx	x	x			x	x	
10	CMM	CMM	CMM x Failsafe Load Required	x	x		x		x	
11	CMM	CMM	CMM x Failed to Start Test xx	x	x		x	x	x	
12	I2C	CMM, MBM	I2C Failure on MBM xx CMM x	x	x	x				x
13	MBM	MBM	MBM Post Failure xx	x	x			x	x	
14	Power Supply	MBM	Power Supply x can't power on	x	x	x				x
15	LAN	MBM	No Peer Communication on LAN	x		x				x
16	UART	MBM	COM Port x Failure	x	x					x
17	MBM/PBM	MBM/PB M	Fail Safe Load Required	x	x					x
18	LAN	MBM/PB M	In Isolated Group on LAN	x	x	x				x
19	Memory	MBM/PB M	Single/Multi bit Ecc Error	x	x			x	x	
20	WDT	MBM/PB M	MBM Watchdog Reset	x	x	x				x
21	WDT	MBM	Watchdog Expired on Partition xx	x	x	x				x
22	IO7	PBM	IO7 xx Drawer not accessible	x	x	x				x
23	Temp	MBM/PB M	Temperature too high xxx Fahr	x	x	x	x			x
24	EEROM	CMM, MBM/PB M	EEROM xx not accessible	x	x					x
25	OCP	MBM	OCP not accessible	x	x					
26	IO7	MBM	EV7 xx not connected to IO7 xx	x		x				x
27	EV7	MBM	Virtual Console at EV7 xx busy	x						x

Error Codes- Fig. 87




Shared Ram Communiaction

Fig. 88




MBM to CMM communication

Fig. 89




Example of MBM forwarding

Fig. 90




Example of CMM forwarding

Fig. 91



CMM com port connection

Fig. 92



The Telnet Session

Fig. 93

### REQUEST FORMAT

Size (dec)	Start (hex)	End (hex)	Bit 7	Bit 6	Bit 5	Bit 4	Bit 3	Bit 2	Bit 1	Bit 0
4	0	3								Originator IP address
4	4	7								Destination IP address
4	8	B								Identifier
2	C	D								Command Code
n	E	n+E								Data (optional)

Request Format

Fig. 94

## **RESPONSE FORMAT**

## Response Format

Fig. 95

Bit 7	Bit 6	Bit 5	Bit 4	Bit 3	Bit 2	Bit 1	Bit 0
Originator IP address (0)							
Destination IP address (FFFFFFFh)							
Identifier (GroupID)							
Command Code							
SM LAN message							

Train Message Header Format

Fig. 96

Command Descriptor	Code
New group	0101h
Accept group offer	0102h
Reject group offer	0103h
Join group	0104h
Probe microprocessor	0105h
I-am-alive	0106h
Report Conflicting Address	0107h
Set Membership Configuration	0108h

*LAN Formation Group*

Fig. 97

Command Descriptor	Code
Full Train Message	0201h
Empty Train Message	0202h

### Reliable Message Group

Fig. 98

Command Descriptor	Code
Get CMM State	0310h
Get MBM Configuration	0321h
Get PBM Configuration	0322h
Get Partition Database	0323h
Distribute Partition Database	0324h
Get System Topology	0330h
Store PCI Slot Info	0331h
Get PCI Slot Info	0332h
Get Own Partition Number	0333h

#### System Discovery Group

Fig. 99

Command Descriptor	Code
Create Partition	0401h
Set Partition Attributes	0402h
Move EV7s to Partition	0403h
Remove EV7s from Partition	0404h
Save Partition Assignment	0405h
Start Partition	0406h
Reset Partition	0407h
Power On Partition	0408h
Power Off Partition	0409h
Halt Partition	040Ah
Add EV7s to Running Partition	040Bh
Delete EV7s from Running Partition	040Ch
Switch Primary EV7	040Dh
Destroy Partition	040Eh
Continue Partition	040Fh
Compute Routing	0410h
Configure RBOX/CBOX	0411h
Set Partition State	0412h
Get State of OCP switches	0413h
OCP Switch Assignment	0414h
Power On/Off	0415h
System Event	0416h
Assign Sub Partitions to Community	0417h <sup>1</sup>
Get Hard Partition Memory Assignments	0418h
Assign Memory Block to Sub Partition	0419h
Assign IO7 to Sub Partition	041Ah
Store Environment Variables	041Bh
Get Environment Variables	041Ch

### Partition Control Group

Fig. 100

---

<sup>1</sup> Communities are to be implemented at a later phase of development.

Command Descriptor	Code
EV7 Reset On/Off	0501h
EV7 Pulsed Reset	0502h
EV7 Halt On/Of	0503h
EV7 Quiesce	0504h
EV7 RBOX/CBOX Config	0505h
Request EV7 Start Test	0506h
Load Image	0507h
Load & Run SRM	0508h

#### EV7 Setup Group

Fig. 101

Command Descriptor	Code
Set Cable Test Signal State	0601h
Get Cable Test Signal State	0602h
Send Cable ID	0603h
Receive Cable ID	0604h
Get MBM IP Cabling	0605h
Get PBM IO Cabling	0606h
Get Cabling Configuration	0607h
Reconfigure Cabling	0608h

#### Cable Test Group

Fig. 102

Command Descriptor	Code
Get Telnet IP Address/Port	0701h
Put Chars from Keyboard to Virtual Cons	0702h

Virtual Console Group

Fig. 103

Command Descriptor	Code
Get Firmware Version	0801h
Upgrade Firmware	0802h
Load Test Version	0803h
Disable Test Version	0804h

#### Firmware Load and Upgrade Group

Fig. 104

Command Descriptor	Code
Get Voltage Readings	0901h
Get Temperature Readings	0902h
Get Fan RPM Readings	0903h
Set Fan RPM Speed	0904h
Set OCP Display Data	0905h
Set Attention Indicator	0906h
Get Switch State	0907h
Get Power Supply State	0908h

#### Environmental Retrieval Group

Fig. 105

Command Descriptor	Code
Get EEROM Data	0A01h
Set EEROM Data	0A02h

FRU Data Group

Fig. 106

Command Descriptor	Code
Error Reporting	0B01h
Get Error Log Count	0B02h
Error Log Clear	0B03h
Get Error Log Entry	0B04h

Error Logging Group

Fig. 107

Command Descriptor	Code
Start OS Watch Dog	0C01h
Keep Alive	0C02h
Stop OS Watchdog	0C03h

OS Watch Dog Timer

Fig. 108

Command Descriptor	Code
Get Base Time	0D01h
Set Base Time	0D02h
Distribute Base Time change	0D03h
Set Partition Delta Time	0D04h
Get Partition Delta Time	0D05h

Date/Time Group

Fig. 109

Command Descriptor	Code
Get Knob	0E01h
Set Knob	0E02h
Unrecognized Response	0E03h
Distribute DHCP Lease Data	0E04h
Read	0E05h
Write	0E06h

#### Miscellaneous Group

Fig. 110

## DECLARATION AND POWER OF ATTORNEY FOR PATENT APPLICATION

As a below-named inventor, I hereby declare that:

My residence, post-office address, and citizenship are as stated below next to my name.

I believe I am an original, first, and joint inventor of the subject matter which is claimed and for which a patent is sought on the invention entitled PARTITION CONFIGURATION VIA SEPARATE MICROPROCESSORS, the specification of which is attached hereto and identified by Cesari and McKenna File No. 15311-2284.

I hereby state that I have reviewed and understand the contents of the above-identified application specification, including the claims, as amended by any amendment specifically referred to herein.

I acknowledge the duty to disclose all information known to me that is material to patentability in accordance with Title 37, Code of Federal Regulations, §1.56.

I hereby claim foreign priority benefits under Title 35, United States Code §119(a)-(d) of any foreign application(s) for patent or inventor's certificate listed below and have also identified below any foreign application for patent or inventor's certificate filed by me on the same subject matter having a filing date before that of the application on which priority is claimed: None.

I hereby claim the benefit under Title 35, United States Code §119(e) of the following U.S. provisional application: None.

I hereby claim the benefit under Title 35, United States Code §120, of the United States Application(s) listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States application in the manner provided by the first paragraph of Title 35, United State Code, §112, I acknowledge the duty to disclose all information that is material to patentability in accordance with Title 37, Code of Federal Regulations, §1.56, and which became available to me between the filing date of the prior application and the national or PCT international filing date of this application: None.

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment or both under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

I hereby appoint Michael E. Attaya, Reg. No. 31,731; Charles J. Barbas, Reg. No. 32,959; Joseph H. Born, Reg. No. 28,283; John L. Capone, Reg. No. 41,656; Robert A. Cesari, Reg. No. 18,381; Yong S. Choi, Reg. No. 43,324; Brian C. Dauphin, Reg. No. 40,983; Steven J. Frank, Reg. No. 33,497; Howard S. Fuhrman, Reg. No. 33,175; Christopher K. Gagne, Reg. No. 36,142; A. Sidney Johnston, Reg. No. 29,548; William A. Loginov, Reg. No. 34,863; John F. McKenna, Reg. No. 20,912; Rama B. Nath, Reg. No. 27,072; Martin J. O'Donnell, Reg. No. 24,204; Thomas C. O'Konski, Reg. No. 26,320; Edwin H. Paul, Reg. No. 31,405; Michael R. Reinemann, Reg. No. 38,280; Rita M. Rooney, Reg. No. 30,585; Heather B. Shapiro, Reg. No. 41,305; Patricia A. Sheehan, Reg. No. 32,301; and Joseph Stecewycz, Reg. No. 34,442, Cesari and McKenna, LLP, 88 Black Falcon Avenue, Boston, Mass. 02210, and Irene Kosturakis, Reg. No. 33,724; Keith Lutsch, Reg. No. 31,851; Barry Blount, Reg. No. 35,069; Sarah T. Harris, Reg. No. 35,891; Joseph Arrambide, Reg. No. 39,589; Richard P. Lange, Reg. No. 27,296; Theodore S. Park, Reg. No. 26,971; and Anirma Rakshpal Gupta, Reg. No. 38,275,, , jointly, and each of them severally, my attorneys and attorney, with full power of substitution, delegation and revocation, to prosecute this application, to make alterations and amendments therein, to receive the patent and to transact all business in the Patent and Trademark Office connected therewith. Please direct all telephone calls to A. Sidney Johnston at. Please address all correspondence to A. Sidney Johnston.

David Golden	Date
Residence:	80 Morin Street Woonsocket, RI 02895
Citizenship	USA
Post Office Address:	Same as Above

---

Dennis Mazur	Date
Residence:	14 Saybrook Road Worcester, MA 01604
Citizenship	USA
Post Office Address:	Same as Above

---

Richard Braken	Date
Residence:	311 Commonwealth Avenue #40 Boston, MA 02215
Citizenship	USA
Post Office Address:	Same as Above